Footnote Number	Content
1	“The science and engineering of making intelligent machines” according to Professor Christopher Manning. https://hai.stanford.edu/sites/default/files/2020-09/AI-Definitions-HAI.pdf
2	These sensors vary from industry to industry, but the objective is to map raw inputs to process and performance parameters, as noted in the case for paper mills by Bain and Co. (https://www.bain.com/insights/mill-of-the-future-paper-and-packaging-report-2023/). For a survey of sensor types within an Industry 4.0 proposal, see Javaid et al 2021 (https://www.sciencedirect.com/science/article/pii/S2666351121000310). For an example sensor design, see Ejeian et al 2019 for a MEMS-based fluid flow rate sensor design. These in turn operate similarly to the Omron D6F series sensor, wherein they convert flow into I2C signal, and then subsequently into Ethernet packets.
3	Decision systems can vary drastically in computational requirements, analytical assumptions, and data requirements. A heuristic system could be a decision tree, which could classify a human based on entities with two eyes and two legs. This assumes inputs would not include primates other than humans. A more sophisticated system could use statistics and conditional probability. If it has two eyes, maybe based on data there is an eighty percent chance for a human to have two legs, and we rely on the classification rate as a tolerance to assume human. We would want to re-examine the data frequently to update our hypothesis, and perhaps when the conditional probability drops to 50/50, we use a different method. A more independent system could be a statistical classifier that learns on the data itself, perhaps via regression analysis. With a broader selection of animal attributes in a training set, it would be able to use linear least squares to minimize the error of selection, given the input attributes. Efficiency can vary depending on the model—classical linear models tend to be sensitive to using many variables in their specification, or when data is non-normally distributed. For these cases, more advanced machine learning can be useful, such as Random Forest regression. Extreme cases with many more classes than examples of data typically are the type where self-attention based Transformer models like GPT are most useful.
4	As with our previous example, there are many factors that can dictate the right model for the circumstance. The general best practice is to follow the scientific method, which tends to favor the principle of parsimony. Simpler models, like the simple classification tree, tend to be easier to analyze, faster to execute for a CPU, and have a much smaller requirement for a working set in RAM. On the other side of the spectrum, large language models can be extremely expensive to maintain. Training a model like GPT-4 from scratch is a multi-million-dollar affair. BloombergGPT (https://arxiv.org/abs/2303.17564) required 1.3 million GPU hours to train. Retraining on new data would require another million GPU hours. Some approaches avoid going to scratch, and instead use new data to fine tune existing models, and that in turn is usually only a fraction of cost. But it’s rare for researchers to know the true difference between a versioned model and merely a fine-tuned model. Given this ambiguity OpenAI researchers typically will train an entirely new model with higher capabilities and more parameters than redo their work. More often, performance requirements for developing models are less complex and expensive than with the large language model. Regression models. When to retrain a model and its acceptable use depends on how much its error and goodness of fit are with a hold-out set, a procedure known as cross validation. Discovered error and goodness of fit in turn dictate how often a model is retrained, how large the model is, and so forth.
5	5 See: https://www.technologyreview.com/2024/06/10/1093417/how-digital-twins-are-helping-scientists/ for more discussion
6	https://www.ascm.org/ascm-insights/the-xyzs-of-inventory-management/ ABC/XYZ, https://ieeexplore.ieee.org/abstract/document/8999262 for forecasting, and https://aws.amazon.com/blogs/publicsector/the-us-air-force-improves-aircraft-readiness-with-ai-and-predictive-maintenance-solutions/ for predictive maintenance
7	An example of support equipment being selective motion control gantry systems ala https://www.assemblymag.com/articles/92437-automation-helps-productivity-take-off-at-boeing or automated picking system ala https://us.blog.kardex-remstar.com/automated-storage-and-retrieval-systems-asrs
8	https://dynamics.microsoft.com/en-us/customer-voice/what-is-the-voice-of-customer

9	Trends are favoring data sovereignty rules in the wake of data leaks. Here are some data sovereignty regimes one should be aware of when running and international operation. General Data Protection Regulation (GDPR). This applies in the EU, and provides rights to consumers over the treatment of their data, such as the right for their data to be forgotten (purged) by the operator. Personal Data Protection Act 2012 (PDPA). This applies in Singapore. Regulates the collection, consent, and resale of data. Personal Information Protection Law of the People's Republic of China (PIPL). This applies in the People's Republic of China. While it gives consumers the right to know and erase data, most significantly it mandates strict regulations over cross-border data exchange. It outlaws all exchange of domestic data to foreign government handlers, even for cases of bind contracts. California Consumer Privacy Act (CCPA). This applies in the US (California). Consumers have the right to view, access, and deny sale of their data. Children's Online Privacy Protection Act (COPPA). This also applies in the US (California) and dictates rules for holding and forgetting data for children under the age of 13.
10	An example, by breaking down a managed IOT service. Many IOT broker services tend to use MQTT (https://mqtt.org/mqtt-specification/), a protocol running on the TCP/IP stack typically on ports 1883 and 8883 for unencrypted and encrypted clients respectively. It is commonly used for remote systems because it offers very strong reliability claims, allowing buffering and retransmission for cases of degraded connectivity, with little effort for the engineer. This client-server infrastructure is often replicated for cloud and equipment provider offerings for managed IOT services. AWS in IoT Core, Siemens in Insights Hub, and Azure IoT Hub are popular options. Each offer SDKs for client development as well as fleet management, over the air updates, analytics and reporting, and so forth. A solutions provider will thus build their own IoT client-server architecture atop this managed service environment. It's important to remember that these examples are all open platforms—it's possible for any company with resources to build their own systems as an alternative to using a SaaS option.
11	One of the most consequential data leaks in history is unfolding as this very article is written (https://www.wired.com/story/epam-snowflake-ticketmaster-breach-shinyhunters/). Snowflake provides a cloud-based data warehousing solution, very popular for predictive analytics operators in the industrial space--both Siemens and McKesson are customers. Recently, they had customer data breached via a subcontractor based in Belarus (EPAM Systems). That subcontractor had detailed access to customer accounts including secret keys, via the issue tracking software Jira. When an EPAM contractor was hacked, that hacker got access to Jira, and thus got free access to the full database of every customer listed in Jira. This was limited to customers who did not have a second factor for that access pattern. But this did include a leak of customer and process data for major corporations. While there is no confirmed case of manufacturing firms affected by this incident, the lack of segregated user access controls, of multifactor authentication were things that the customer could have controlled to prevent global access. Further, EPAM itself could have used a secrets management system to avoid disclosure within the clear in Jira, to reduce the liability for its customers. A very prescient story that could easily have affected a manufacturing company, and could be managed or prevented with defense in depth.
12	While the textual reference concerns CVE 2023-39659, in total security personnel should consider these four categories of compromise that can affect a decision system.
13	https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1800-10.pdf and https://www.nist.gov/system/files/documents/2017/06/12/manufacturing_profile_29apr2016.pdf for the guidance and the profile respectively
14	In particular, there are two kind of solution categories that apply here - cloud access security brokers (CASB) and secure access service edge (SASE). CASB can account for solutions like Skyhigh, which intermediates access between users and a cloud provider, spotting potential intrusion on enterprise assets. SASE is secure access service edge, which for this discussion is secure access to OT resources but could include an entire enterprise. C.f. TxOne EdgeFire and ZScaler ZPA
15	https://www.mitre.org/sites/default/files/2023-01/PR-22-2824-Crown-Jewels-for-Industrial-Control-Systems.pdf

16	Canary token is ala Thinkst Canary, a token that is like real tokens of a given schema, except there are online scanners designed to pick up the moment the Canary is used, and to include metadata from the attempt. It can be put into vulnerable systems to anticipate lateral movement by bad actors. Take the Snowflake case. Should no tokens be disclosed in Jira, then the adversaries would not have been able to access customer accounts without additional probing. If the Canary were used defensively, then the contractors might have been able to minimize time to alert for probing. It's worth noting there are also threats to exposure for a model itself in some cases. Nevo et al 2024 (https://www.rand.org/pubs/research_reports/RRA2849-1.html) discuss cases of model weight exfiltration. This applies to cases where your organization has created its own model which contains sensitive information or otherwise contains sensitive intellectual property. These cases are typified by public or expansive access to the model serving or training infrastructure, so the use of canary tokens and least privilege access patterns help mitigate this risk as well.
17	Human in the loop is a control system pattern that uses predictive systems, but subordinate to the judgement of one or more human inputs. It is frequently used in simulation studies (e.g. Consiglio 2010) to create assistive systems, owing to humans' ability to add high-quality labeled data to the predictive system and the automated system in managing normal conditions. Bhattarchaya et al 2023 (https://www.mdpi.com/2079-8954/11/1/35) create an example of this on the factory floor, where the immediate product of interest was human machine interfaces that could use predictive systems with human input for creating rich context, such as providing shop floor conditions when evaluating the performance of an individual machine.
18	Statistical control charts are used to determine the range of valid, non-spurious configurations for systems within a control context. https://deming.org/a-beginners-guide-to-control-charts/ - they give guidance for detecting anomalous behavior that is actionable without prior data. Not all processes fail in a way that is consistent with control chart rules, so it is advisable to tune each metric's outlying conditions using a historical dataset with statistical significance carefully if they want to expand on the control chart methodology.