COMPREHENSIVE MACHINE-AUGMENTED ASSURANCE

How machine learning can improve assurance and make smart people smarter

Augmented decision-making tools can improve both the efficiency and effectiveness of expert-led assurance activities. In turn, these tools can deliver significant business benefits, such as by avoiding delays in acceptance of new systems, reducing costs, and de-risking the system that is being assured. In this Viewpoint, we share ways to utilize such technology to the greater benefit.

AUTHORS

Clive Adams

David Boulton

James Beck

Michael Eiden

Philippe Monnot

Armand Rotaru

BREAKING THE TRADITIONAL ASSURANCE MOLD

Traditional assurance activities require expert engineers to undertake a detailed review of extensive volumes of documentation, which is effort-intensive, time-consuming, and costly. However, carefully applied augmented decision-making tools can replace some of the laborious manual checks normally required (see Figure 1). The aim is not to replace human experts but to improve the quality and efficiency of decisions they make.

In this Viewpoint, we demonstrate how augmented decision-making tools can be used in practice with independent safety assessments (ISAs). However, the tools and techniques are applicable to a wide variety of other use cases and scenarios.
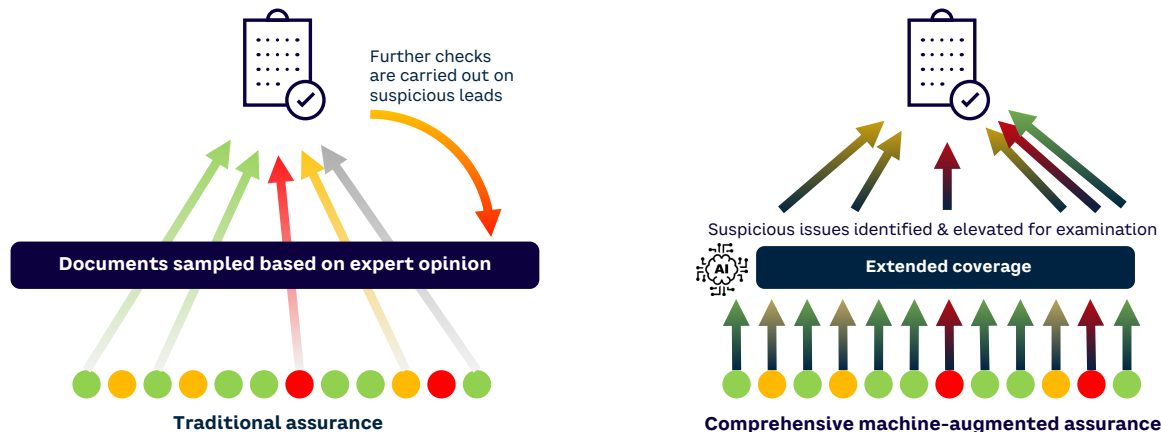
WHY USE MACHINE LEARNING?

ISA involves making an expert-led judgment on: (1) whether the requirements that have been placed on a system are sufficient to ensure its safety and (2) whether the system has met those requirements. ISA therefore provides increased assurance that a system can be deployed and operated safely. It also typically requires use of expert engineers with many years' experience in the associated industrial domain.

Expert human knowledge and insights are notoriously difficult to encode in conventional computer algorithms. However, machine learning (ML) offers a solution to this problem. An ML tool can be trained to make inferences about the data supplied to it based on inferences that it has learned from past "training data." Such a tool can then make independent judgments on new data, without human intervention.

IT'S ALL ABOUT RELATIONSHIPS

Much of an ISA involves considering the correctness and completeness of the relationships between different data sets (e.g., whether system design requirements fulfill safety requirements or whether subsystem design requirements correctly implement system design requirements). Assessments sample check these relationships. An expert assessor chooses a level of sampling that allows him or her to reach a reasonable conclusion on the quality of the data concerned. The greater the degree of integrity a system requires, the greater the breadth and depth of the sample checking that is needed.

Figure 1. Traditional vs. machine-augmented assurance

Further checks are carried out on suspicious leads

Documents sampled based on expert opinion

Traditional assurance

Suspicious issues identified & elevated for examination

Extended coverage

Comprehensive machine-augmented assurance

Source: Arthur D. Little
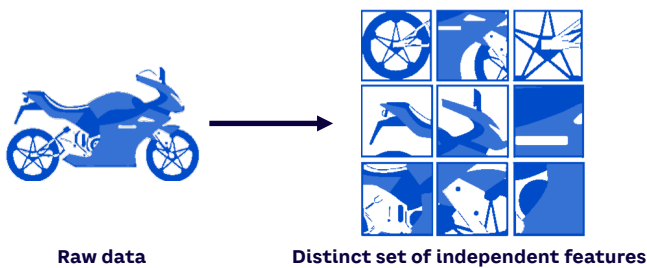
MANAGING LARGE VOLUMES OF DOCUMENTATION

System development can involve huge volumes of documentation (data sets). One of the first challenges is for an ML model to ingest such documentation in a manner that allows for easy reasoning. **Natural language processing** (NLP) techniques are therefore used to extract the relevant data contained within each document, as well as to capture the relationships between them.

Another key challenge is capturing the meaning of each data item to improve the ability to reason about it. **Feature engineering** is therefore used to transform the raw data into attributes that better represent the data (see an example in Figure 2).

ONE OF THE FIRST CHALLENGES IS FOR AN ML MODEL TO INGEST DOCUMENTATION IN A MANNER THAT ALLOWS FOR EASY REASONING
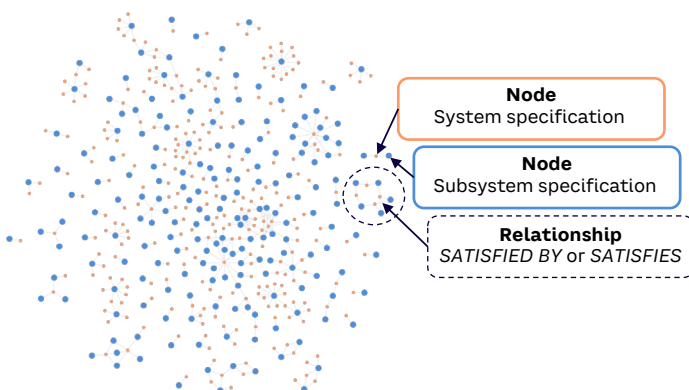
The extracted features and relationships can then be represented semantically using structured **knowledge graphs** (as illustrated in Figure 3). These transform the unstructured data contained in the source documents into a data set that clearly and precisely embodies the many contextual relationships between the different data items.

Figure 2. Feature engineering

Raw data　　　**Distinct set of independent features**

Source: Arthur D. Little

Figure 3. Representing data sets using knowledge graphs

Node
System specification

Node
Subsystem specification

Relationship
SATISFIED BY or *SATISFIES*

Source: Arthur D. Little

LEARNING FROM EXPERIENCE

Once the large volume of source data has been structured into knowledge graphs, iterative processes are used to develop and train an ML model to implement given use cases. As noted earlier, in the case of ISA, a typical use case is determining the correctness and completeness of requirement traceability from one level of requirements to another (e.g., from system-level requirements to software-level requirements). A combination of different classes of features can be used to train the ML model for such a use case, as illustrated in Figure 4.

In this way, the model can emulate the pattern matching that is intuitively performed by expert human engineers. This includes, for example, evaluating the similarity of low-level data (e.g., functions, variables, constants, expressions) included in the documents. The structured representation of data in knowledge graphs facilitates the ML model to achieve such pattern matching.
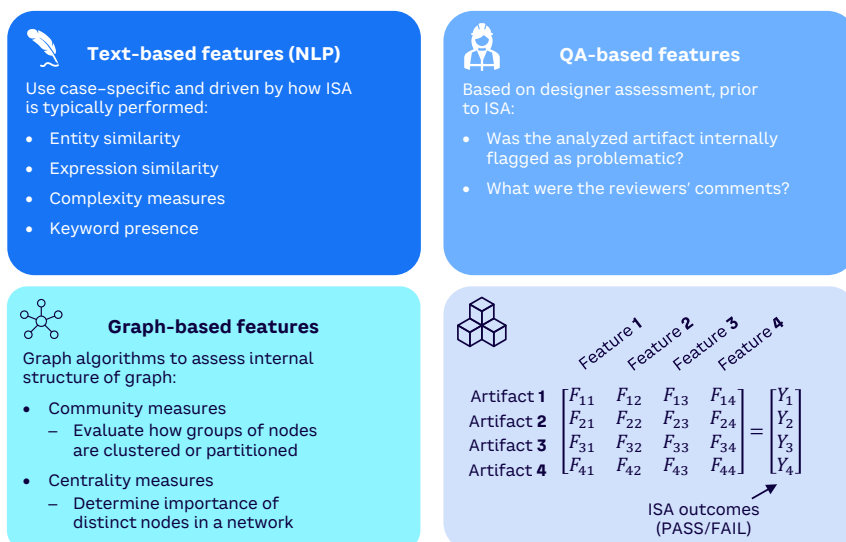
The model learns from "past experience" by using input from expert engineers. For example, such input may reveal which relationships between system and software requirements experts have previously considered correct and which they have not.

TUNING THE ML MODEL

The performance of the ML model is improved in an iterative process, as shown in Figure 5. In particular, using human experts' insight, the model can be trained on an increasing number of training data sets. Such data sets are those in which experts have already made correct inferences on the relationships between the data that they contain. As the machine's training improves, it is increasingly able to emulate the subjective decisions and insights that the expert human engineers employ (decisions and insights that such experts may find difficult to articulate and/or which they may make subconsciously).

> **AS THE MACHINE'S TRAINING IMPROVES, IT IS INCREASINGLY ABLE TO EMULATE THE SUBJECTIVE DECISIONS AND INSIGHTS THAT THE EXPERT HUMAN ENGINEERS EMPLOY**
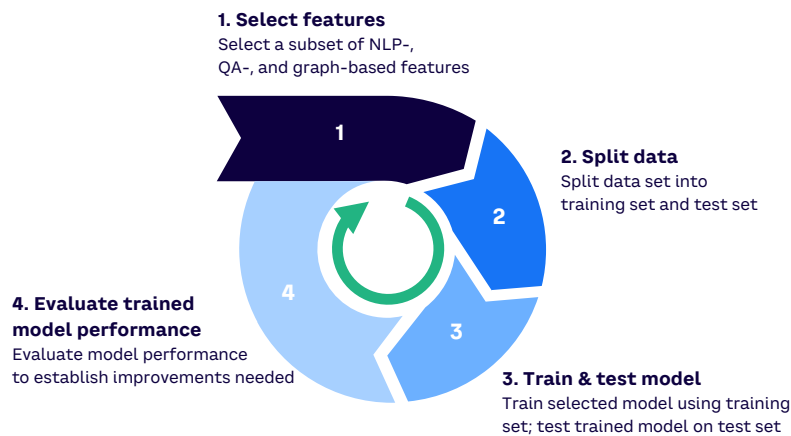
Figure 4. Training ML models using different classes of features

Text-based features (NLP)
Use case–specific and driven by how ISA is typically performed:
- Entity similarity
- Expression similarity
- Complexity measures
- Keyword presence

QA-based features
Based on designer assessment, prior to ISA:
- Was the analyzed artifact internally flagged as problematic?
- What were the reviewers' comments?

Graph-based features
Graph algorithms to assess internal structure of graph:
- Community measures
  - Evaluate how groups of nodes are clustered or partitioned
- Centrality measures
  - Determine importance of distinct nodes in a network

$$\begin{array}{c} \text{Artifact } \mathbf{1} \\ \text{Artifact } \mathbf{2} \\ \text{Artifact } \mathbf{3} \\ \text{Artifact } \mathbf{4} \end{array} \begin{bmatrix} F_{11} & F_{12} & F_{13} & F_{14} \\ F_{21} & F_{22} & F_{23} & F_{24} \\ F_{31} & F_{32} & F_{33} & F_{34} \\ F_{41} & F_{42} & F_{43} & F_{44} \end{bmatrix} = \begin{bmatrix} Y_1 \\ Y_2 \\ Y_3 \\ Y_4 \end{bmatrix}$$

(column headers: Feature **1**, Feature **2**, Feature **3**, Feature **4**)

ISA outcomes (PASS/FAIL)

Source: Arthur D. Little

Figure 5. Iterative tuning of performance of machine model

1. Select features
Select a subset of NLP-,
QA-, and graph-based features

2. Split data
Split data set into
training set and test set

**4. Evaluate trained
model performance**
Evaluate model performance
to establish improvements needed

3. Train & test model
Train selected model using training
set; test trained model on test set

Source: Arthur D. Little

Industry-standard artificial intelligence (AI) metrics, precision and recall, can also be used to tailor the nature of the results provided by the model. For example, precision and recall are inherently linked, in the form of a trade-off. An appropriate trade-off policy can therefore be chosen depending on the nature of the assessment required:

— **Conservative policy (high precision, low recall).** If the model says the relationship is a FAIL, you can almost be sure that it is an actual FAIL. But you may miss many (less certain) actual FAILs. In this case, expert assessors have good confidence in the documentation and want to focus on identifying the likely small set of errors.

— **Liberal policy (low precision, high recall).** The model will wrongly label a lot of relationships as FAILs, but it will also detect most of the actual FAILs. In this case, expert assessors have lower confidence in the documentation and need to "cast the net wide" to ensure they catch as many errors as possible.

AN APPROPRIATE TRADE-OFF POLICY CAN BE CHOSEN DEPENDING ON THE NATURE OF THE ASSESSMENT REQUIRED

The ML tool can provide the expert assessor with a prioritized list of relationships for examination. By choosing the highest-priority relationships in the list (those that are most likely to be true FAILs), those areas of the development that are incorrect are more readily identified. This approach also has the advantage of identifying areas of the development that may be most prone to failure (by identifying anomalies in a common area), where more expert investigation may be warranted. Such an approach may prove more efficient than selecting a sample of relationships for assessment based only on the expert's intuition.

OTHER BENEFITS

Augmented decision-making tools provide additional benefits as well, including:

- **Faster outcomes.** Replacing some of the more effort-intensive and labor-intensive manual checks of compliance with processes guarantees faster outcomes, while still providing the same (or potentially higher) levels of assurance.

- **High scalability.** Persistent and highly reusable knowledge graphs can be readily expanded with additional documents and use cases.

- **Timely delivery and de-risking.** The efficiencies that can be generated increase the likelihood of on-time delivery of assurance projects, thereby reducing the chance of cost overruns and the need for use of contingencies. Increased assurance can also de-risk the system that is being assured.

- **Expert integration.** The more the machine can incorporate the insights and tacit knowledge held by expert human engineers, the better the potential efficiency and assurance gains. Use of knowledge graphs is a major enabler in this regard. As expert human engineers reason about the data sets by navigating the knowledge graph, the machine can learn the reasoning of the experts.

- **Applicable to a wide range of use cases.** One of the most exciting aspects of this technology is its applicability to many different use cases, not just ISA. Requirements verification, a simple extension of ISA, is one such example. The technology has been applied to the verification of compliance of public service providers' policy and procedures to regulatory requirements, with promising results.

CONCLUSION

FUTURE ASSURANCE

MACHINE LEARNING CAN IMPROVE THE EFFECTIVENESS AND EFFICIENCY OF HUMAN EXPERT-LED ASSURANCE

Augmented decision-making tools and techniques are applicable to a wide variety of use cases and scenarios, offering opportunities for improved quality and efficiency of human expert-led decisions. By helping to avoid delays in integrating new systems, reducing costs, and de-risking the system that is being assured, machine learning can improve the effectiveness and efficiency of human expert-led assurance, offering several benefits for many businesses, including:

1   **Wide applicability.** The tools and techniques employed have broad applicability. Any activity where human experts must reason about the relationships involved in large data sets is a candidate for the tool.

2   **Making smart people smarter.** The aim is not to replace human experts but to augment their decisions with the assistance of ML tools and enhancing the quality and efficiency of the decisions they make.

3   **Application across a wide range of industrial domains.** The technology is domain-agnostic and, once the rules have been learned, is independent of the industry.

ARTHUR·D·LITTLE

Arthur D. Little has been at the forefront of innovation since 1886. We are an acknowledged thought leader in linking strategy, innovation and transformation in technology-intensive and converging industries. We navigate our clients through changing business ecosystems to uncover new growth opportunities. We enable our clients to build innovation capabilities and transform their organizations.

Our consultants have strong practical industry experience combined with excellent knowledge of key trends and dynamics. ADL is present in the most important business centers around the world. We are proud to serve most of the Fortune 1000 companies, in addition to other leading firms and public sector organizations.

For further information, please visit www.adlittle.com.