Amstelveen

Model Assurance

Maintaining trust in machine-based decision making

WHITE PAPER BY **ED LITTLE**

Computer-based Models are commonly used to gather information and make automated decisions. When they go wrong the implications can be significant.

In this article we discuss what Models are, how they are used, and how controls, regular testing and assurance activities can be used to ensure Models continue to operate effectively.

Organisations that implement Model Assurance programs can reduce the risk associated with developing and operating Models. The resulting safeguards and increased transparency can increase management confidence in the operation of Models and may open new opportunities previously considered too risky.

Background

What are Models, and why do we use them?

Models are simulations which are created to make predictions about outcomes or behaviours. These simulations may occur by the execution of a process or a computer-based algorithm that mirrors or performs a task. Models have been used for centuries by financial institutions to assess risk, score credit or assess the likelihood of an outcome.

With the increase in data available to organisations, the use of Models has also increased. Everyday Model applications include:

- Forecasting the weather;
- Filtering the spam out of email inboxes;
- Helping us navigate while we are driving, or even driving our cars for us;
- Recommending new music to listen to; and
- Unlocking our phones with facial recognition or a fingerprint reader.

Although Models are powerful tools with widespread applications, there are some downsides that need to be managed.

Risks in the use of Models

There have been some very high-profile Model failures in the media over the last decade.

In 2016 it was reported that the Australian Federal Government had incorrectly pursued welfare recipients for debt collections using an autonomous system in what became known as ‘robodebt’ and ended in an investigation from the Commonwealth ombudsman and a class action which was settled for $1.2bn in November 2020.

In the United States, issues related to Model bias against segments of the population are well documented. One high profile failure was Apple’s initial implementation of facial recognition software ‘FaceID’. Siblings of iPhone owners were able to unlock the owner’s device without first adding their face to the safe list.

We are unlikely to prevent all Model failures, but there are some steps we can take to help reduce the risk associated with operating and relying on a Model.

Model Assurance

Model Assurance has two primary components; validating the outputs of the Model (Model Validation), and verifying that controls are in place to maintain the integrity of these outputs (Model Governance).

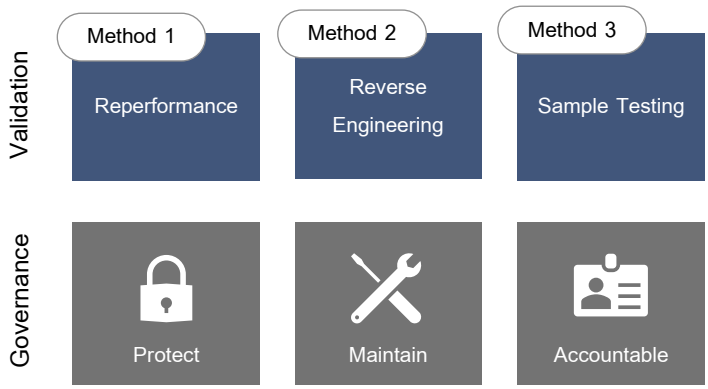

Model Validation

The most important part of a Model Assurance engagement is to confirm that the Model accurately produces an expected outcome based on a known set of inputs. There are three main approaches to validation which relate to the level of assurance required and time available.

Validation Method 1: Reperformance

The most comprehensive and effective means of testing the design of a Model is to build a replica based on the known, documented requirements and testing the performance of both Models against a known dataset.

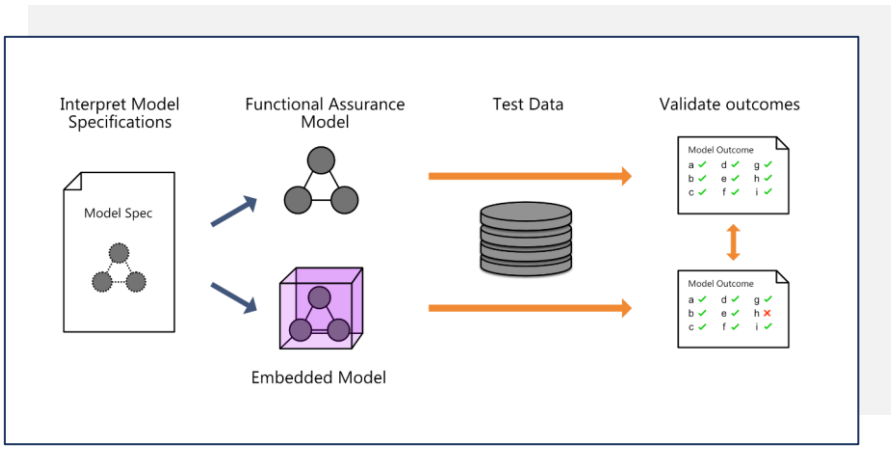

Validation Method 2: Reverse Engineering

Where Model requirements have been poorly documented or maintained, or the Model is a “black box” (a system where the internal workings cannot be observed), it may be possible to reverse engineer the Model and recreate the requirements. These requirements can then be validated with management to ensure the embedded model is functioning as expected.

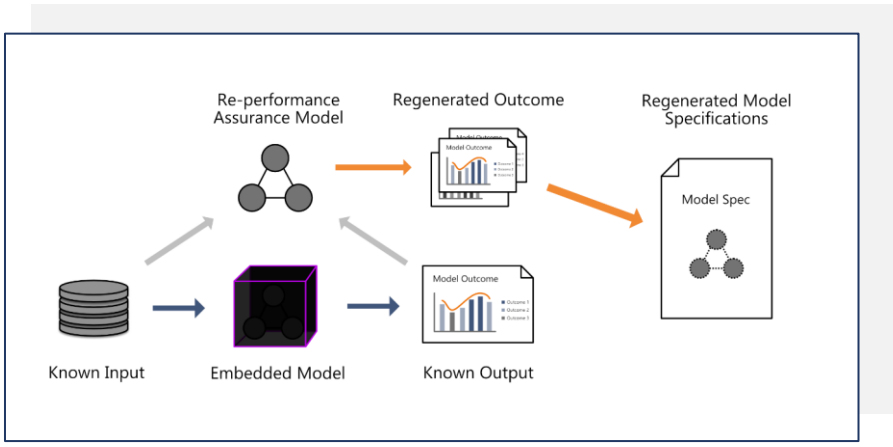

Validation Method 3: Sample Testing

In sample testing, Models are provided with a sample dataset and the output is validated against an expected result. Sample testing is used in situations where a lower standard of assurance is required or there is a short time in which to conduct the testing.

Despite the less comprehensive nature of this testing method, sample testing can be integrated into software testing practices or an Internal Audit program to allow for regular testing of a Model’s performance.

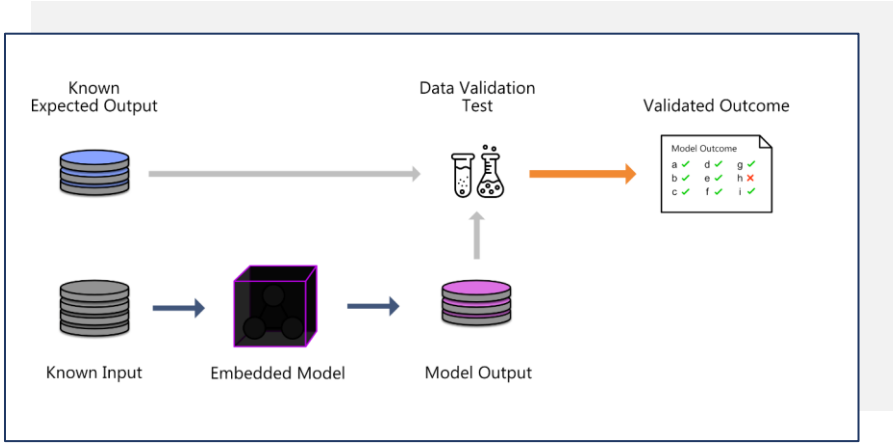

Model Governance

Organisations that rely on Models must establish sound governance principles that ensure that Models remain safe, accurate and supported.

There are three principles relevant to the Governance of Models:

1. Protect your Models.
2. Maintain your Models.
3. Assign an owner to each Model.

Protect your Models

Models may have access to sensitive information, so it is important to classify the data ingested and produced by a Model. Data classification should be aligned to the organisation's information security policy and subject to the relevant controls.

Access to the Model should be restricted to authorised individuals and additional steps taken to identify any changes made to the Model. This can be challenging when Models are stored in spreadsheets. Most spreadsheet software recalculates the state of the sheets when the application is opened or a cell is updated which can obscure other changes made to the document.

Protect the availability of key Models, an outage could disrupt other systems that rely on them or prevent business operations from functioning correctly.

Where access to a Model cannot be restricted, strong version control and change detection controls must be implemented, monitored, and reported on. With the increasing adoption of integrated deployment and release pipelines, there are many tools available to support the effective protection of models.

Maintain your Models

The upfront cost and effort associated with building and implementing a Model can be relatively small, however they can be very difficult and expensive to maintain over time, potentially incurring substantial technical debt[1]. This is because even small changes can have the effect of changing the entire Model. System upgrades or migrations often surface failures in the Model maintenance process which puts the organisation at risk if the model cannot be accurately replicated in the new system.

Models should have a maintenance cycle like any other piece of critical software, supported by a team with the expertise to maintain, test, and update the Model in line with business requirements.

Organisations can easily be put at risk when the staff member who developed the Model leaves and appropriate support structures are not in place.

---

[1]For more on technical debt in machine learning, see Machine Learning: The High Interest Credit Card of Technical Debt D. Sculley et. all — Google research publication, NIPS 2014 workshop.

Assign an owner to each Model

It is common to have 'business' and 'system/technical' owners for production systems who share the responsibilities for designing, operating, and maintaining a system. Equally, assigning an owner to a Model and making them accountable is a key enabler for the effective management and governance of Models.

The main objectives of the person (or persons) accountable for the Model should include:

- **The operation and availability of the Model**
  Ensuring the Model is available when it is required and operating within expected boundaries (including Model performance measures and resource consumption)
- **The security and protection of the Model**
  Ensuring there are adequate protections in place for the Model and the data flowing in and out.
- **The development and maintenance of the Model**
  Ensuring there is sufficient investment in the ongoing development of the Model so business objectives continue to be met.
- **Monitoring Model performance and operation**
  Ensuring the Model is operating within pre-determined boundaries including query response times and data and decision output quality.

The accountable party is ultimately responsible for the enabling the effective design and governance of the Model, thereby ensuring the Model can be certified as fit-for-purpose and operating within expected bounds.

Certifying Models

A Model Assurance review should test the controls and processes that support a Model and certify or attest that the Model is fit-for-purpose and operating within an agreed set of boundaries. The certification can be used to provide assurance to management, the Board, and customers that Models within the organisation are being managed effectively.

To achieve a successful certification result, both aspects of Model Assurance must function effectively together. A strong design cannot be maintained and will deteriorate without strong governance. Committing resources into strong governance should not be attempted without first validating that the design is sound, attempting to do otherwise results in an ineffective use of time, money, and effort.

Conclusion

Models are used to support many aspects of organisational operations and decision-making processes. Organisations should conduct Model Assurance to build trust in key models, which should include the validation of model outputs and assessments of model governance.

Ed Little

Solution Lead

Phone: +61 420 944 732

Email: elittle@amstelveen.com

Ed supports Amstelveen's clients in Model Assurance, Cyber Security, Reporting, Data Analytics and Risk Management engagements.

Amstelveen

Amstelveen's team of professionals work on major technology and business change projects, and enhance capability in risk management, internal audit and corporate governance functions. We provide specialist support to the largest public and private organisations in Australia and New Zealand.