MINISTRY OF SOCIAL DEVELOPMENT
TE MANATŪ WHAKAHIATO ORA

User Guide

Model Development Lifecycle

MSD Reference: A13094445

October 2021

Contents

Introduction

In Aotearoa, the operational use of algorithms to support decision-making is common within government departments[1]. As departments continue to use algorithms alongside other emerging uses of data to support how they operate, they continue to learn about the potential benefits, but also risks and harms and how to best manage these.

The Model Development Lifecycle (MDL) framework is made up of three guides. These guides support organisations to develop operational algorithms to support decision-making in a way that maximises benefits while appropriately managing risks and harms.

These guides have been developed in partnership with the Ministry of Social Development and Nicholson Consulting.

Guides that make up the Model Development Lifecycle

The MDL is made up of the following guides:

1. User Guide (this document)
2. Data Science Guide for Operations (Data Science Guide)
3. Governance Guide

To support the governance framework for operational algorithms there should be a documented sign-off process with approval from each of the following:

- L2 Analytics Owner
- L2 Business Owner
- L2 Communication Owner
- L2 IT Owner
- L2 Analytics Owner - Ethics
- Technical Advisory Group review outcome
- L3 Operational Algorithm Owner (authorising go-live)

The *Data Science Guide* is targeted towards data scientists while the *Governance Guide* and other supporting materials are targeted towards project managers and coordinators.

[1] https://www.data.govt.nz/assets/Uploads/Algorithm-Assessment-Report-Oct-2018.pdf

Required roles for the Model Development Lifecycle

Overview

To ensure operational algorithm products are built in a way that maximises benefits while appropriately managing risks and harms, accountability through clearly defined roles and responsibilities is crucial.

This section details each role and provides guidance on the skills and responsibilities of each role. This should be used when assigning project members to roles in the development of operational algorithms.

The roles used in the MDL can be assigned to multiple staff.

Roles

Data Scientist

For the purposes of the MDL, data scientist is a broad term that encompasses statisticians, analysts, modelers and many other similar job titles. It is envisaged that it would not be an individual data scientist working on an operational algorithm, but rather a team of people with a range of skills. Within that team there should be a technical lead who oversees the work and advises the analytics owner on whether they can sign off the work.

The role of the data scientist includes the core technical work of building and documenting the model. However, this role appears in every section of the MDL because they must also interface with all the other roles to allow those other roles to complete their part of the product development and deployment. The success of these interactions will ultimately determine the success of the algorithm.

Business Owner

The business owner would typically be a regional or national manager who has extensive experience with frontline staff and business problems. They will bring with them subject matter experts (SMEs) who can provide insights about the data and the real-world decision-making process. In many cases it will be the SMEs who assist the data scientist and other roles to understand the business problem. However, the business owner will be a key decision maker in several areas and is responsible for the sign-off of the business specific parts of the product development.

Analytics Owner

The analytics owner has a key role in ensuring the quality of the analytics component of the product development. They need to understand the work of the data scientists to a sufficient level of detail to be able to sign off the completed model. They will often be advised by the technical lead in this role but need to ask critical questions of the data scientist.

The analytics owner also needs to decide if external reviews are required, based on the recommendation of the Technical Advisory Group, as well as the scope of the review.

Privacy and Ethics Specialist

The privacy and ethics specialist has a critical role in understanding which of the available tools and frameworks is appropriate to use at each point in the product development. Organisations will have their own, more or less, formal Privacy, Human Rights and Ethics

(PHRaE) framework, that in turn may refer to external guidelines and practices. For the bulk of the product development, the PHRaE will be the key tool. However, the 'idea definition' phase only requires a small initial assessment, which may not utilise the PHRaE.

When completing the PHRaE the privacy and ethics specialists need to ensure that the development team can clearly articulate the goal and benefits of the project. They also need to ask critical questions of the team throughout the development process.

Change Manager

The change manager's role will depend on the size of the change involved and how the algorithm affects the experience of staff, service providers and service users. It will also depend on the goal of the product. For example, if the goal is efficiency gains then the change manager will need to work with staff quite extensively. However, if the goal is more consistent decisions, then the change manager will have a much smaller role.

Deployment Manager

The deployment manager will be responsible for implementing the algorithm. The size of their role will depend on the degree of integration with the core system. For close integration with the core system the deployment manager will have several staff working with them. The deployment manager needs to ensure the implementation is thoroughly tested and documented and that the appropriate version controls are in place to allow the algorithm to be updated. Many of these tasks are not familiar to the data scientist so the deployment manager needs to bring the rigor and structure that is standard in IT implementation.

The deployment manager also needs to sign off the implementation, ensuring that the algorithm works in practice and has the appropriate support in place.

Table 1: Roles associated with each section of the Data Science Guide.

Section	Role						
	Data Scientist	Business Owner	Analytics Owner	Privacy and ethics specialist	Change manager	Deployment manager	Communication specialist
Idea formulation, selection and planning	☑	☑		☑			
Model selection and optimisation	☑	☑	☑				
Data preparation	☑	☑	☑				
Fairness and bias	☑	☑	☑	☑			
Model and process maintenance	☑	☑	☑				
External review	☑		☑				
Model and process design	☑	☑					
Consultation and co-design	☑	☑					
Change management	☑				☑		
Leveraging tools and processes	☑					☑	
Transparency and communications	☑						☑
Privacy and ethics	☑	☑	☑	☑			
Risk management	☑		☑				

Required roles for the model governance

Overview

There are five sign-off roles within the *Governance Guide*. Who occupies these roles will depend on the benefits and risks associated with the algorithm. For example, for high risk/benefit algorithms the L3 sign-off may be done by the Chief Executive (CE). On the other hand, for lower risk models the Deputy Chief Executive (DCE) or General Manager (GM) may do the L3 sign-off.

As a general principle, the higher the risks and benefits associated with the algorithm, the higher the positions of each of the people filling the roles detailed below will be.

The sign-off documentation will be filled in by key roles in each of the areas. These are detailed in the *Data Science Guide* and summarised here.

Roles

Table 2: Sign-off roles and responsibilities.

Role	Who would typically fill this role	Sign-off responsibility	Sections from the Data Science Guide covered
L2 analytics owner	Team manager or GM of analytics	L2 analytics owner sign-off (filled in by the data scientist)	– Model selection and optimisation – Data preparation – Fairness and bias – Model and process maintenance – External review
		L2 analytics owner – ethics sign-off (filled in by the data scientist)	– Privacy and ethics
L2 IT owner	Team manager or GM of the team that implemented the algorithm	L2 IT owner sign-off (filled in by the deployment manager)	– Leveraging tools and processing
L2 communications owner	Team manager or GM of the communications part of the organisation	L2 communications owner sign-off (filled in by the communications specialist)	– Transparency and communications
L2 business owner	GM of the business unit most affected by the algorithm or a dedicated business owner for the algorithm	L2 business owner sign-off (filled in by the change manager and the data scientist)	– Model and process design – Consultation and co-design – Change management
L3 operational owner(s)	The DCE of analytics and DCE of Service Delivery. The CE for high risk/benefit algorithms.	L3 Operational algorithm go-live sign-off (filled in by the data scientist and analytics owner)	Overall sign-off

Using the Model Development Lifecycle to update an existing operational model

The *Data Science Guide* and *Governance Guide* address the development, deployment and maintenance of a new model. However, most government organisations have some existing models in operation. These will often need a significant refresh that goes beyond standard maintenance.

A significant refresh is more complicated than simply changing the numbers in the equations because of changes in the environment they are being used in. These include:

- changes in the public's expectations and developments in the understanding of how operational algorithms should be used (eg the expectations of transparency have increased dramatically over the last five years)
- changes in the business process or how the business wants to use an algorithm (eg the business may have relied on paper forms in the past but has now moved to digital platforms)
- changes in the data and the patterns observed (eg the quality of the data collected may have increased once it was used as an input to the algorithm)
- changes in the method or technology used to deploy the model into operation (eg the use of cloud-based solutions).

Table 3 details how relevant each part of the *Data Science Guide* and *Governance Guide* are likely to be to a model refresh and why.

Table 3: Relevance of each of the sections in the Governance and *Data Science guides to a model refresh.*

Section	Percentage relevance	Explanation
Governance	80%	Governance of operational algorithms is one of the fastest areas of evolution in data science. It is now developing a coherent structure that allows the benefits and risks to be more effectively managed. The refreshed model will need to be managed through this new structure, which may differ greatly from the original governance approach.
Idea formulation, selection and planning	40%	Only a limited version of the idea formulation and selection is required since this will primarily amount to refining or confirming the original problem and solution. However, planning will still be an important component.
Model selection and optimisation	20%	In many cases, the same variables and type of model will be used for a refresh. This will require little in the way of model selection and some optimisation.
Data preparation	20%	A new sample of data will be needed. However, we can expect almost all of the variables to be the same and to have similar relationships to the target variable. The data scientist will still need to verify existing patterns

		hold, similar missing values exist and that there are no changes in the meaningfulness of any proxy variables.
Fairness and bias	80%	There have been a lot of changes in public expectation and understanding this area. Most refreshed algorithms will need to be thoroughly tested for fairness.
Model and process maintenance	80%	The appreciation of the importance of model maintenance has increased greatly as organisations have grown used to implementing operational algorithms. The refreshed version will likely need a new monitoring and maintenance plan to ensure that it will continue working in the long-term.
External review	30%	External review is likely to be needed only if new areas of concern are identified or if significant changes have occurred in the environment. For example, ethical considerations have become much more prominent than they used to.
Model and process design	50%	How much work is required in this area will depend on how successfully the model and process are currently integrated. If they are working well together then little effort will be required. However, any significant changes in the model may require changes to the process that goes with it.
Consultation and co-design	70%	There has been a growing appreciation of the importance and usefulness of consultation and co-design. This is both in terms of following a good process and in terms of getting good results, so a refresh is likely to require significant work in this area.
Change management	20%	There is likely to be little new work in this area unless the model and process design changes significantly.
Leveraging tools and processes	60%	The tools and methods used by IT tend to change rapidly so the refresh will need to include adapting to these. Data scientists have grown more accustomed to using version control and IT project planning methodologies, but these are likely to have been only sparingly used in the past.
Transparency and communications	70%	Expectations of transparency have grown rapidly over the last five years. Some existing models will have good transparency and clear communications. However, it is likely that most algorithms will need significant work in this area. Even for models with good historic communications, new communications will need to be developed for the updated version.
Privacy and ethics	60%	The privacy and ethics of operational algorithms is a rapidly evolving area. It now incorporates new areas, such as te ao Māori opportunities and risks, and even existing areas have changed. As such, the PHRaE will need to be completed again and new issues will need to be addressed. Having said that, it is unlikely that the existing algorithms will have been implemented without any consideration of privacy and ethics so there will at least be a starting point in place.

Glossary of terms

Model is an analytical process that interprets or evaluates information to solve a problem. This includes statistical, machine learning or artificial intelligence models as well as simple business rules.

Operational model is a model that informs or makes operational decisions.

Operational algorithm is the combination of an operational model and business process that integrates the decisions with the business workflow.

Product development should be taken to include the development, deployment, documentation and testing of the algorithm.

Operational algorithm project is a project where the purpose is to develop, deploy, and maintain operational algorithms products.

PHRaE is the Privacy, Human Rights and Ethics framework used to manage these issues and risks within a project. Larger organisations may have a formal, centrally managed, process. A more ad hoc approach, based around guidelines from external sources, may be more suitable if resources are limited.

Sign-off is approval of a deliverable by an approved person or role to progress to the next project stage.

Sign-off sheets are part of the MDL governance framework and are practical tools to facilitate structured discussions when approving deliverables.

Fairness relates to the system – the overall treatment of each group by the model and the business process, including the influence of the underlying prevalence in the populations.

Bias focuses on the model part alone and the probabilities it generates.

Protected variables are variables such as ethnicity, gender and age. These are the factors that are prohibited, either morally or by law.

Potential Final Model (PFM) is a model that has been selected as a candidate for the final model and is being subject to testing for fairness, accuracy, ease of implementation etc.

Subject matter experts (SMEs) are key staff from within the business who have a strong understanding of how the process works in practice. They will also know the common challenges that staff, service providers and service users face.