



Figure 11: The iterative process of modelling and decision making. † denotes security touchpoints.

Machine Learning & AI Best Practices

The iterative process of modelling & decision making

The ML process of modelling & decision making (see Figure 11) was first developed by Dr. Andreas Wiegand whilst Chief Scientist at amazon.com in the late 90s. The process has since been evolved by Dr. Andrew Ng (Director Emeritus, Stanford AI Lab) and Dr. Sebastien Haneuse, University of Washington.

1. **Measure:** This is where one captures engagement data and actions such as sales, microconversions, signals, telemetry & model performance.
2. **Describe:** This is where one (*I*) formulates the scientific question
 - understanding the association between two variables
 - prediction of some future event
 - output is a loose framing of the problem domain and held out dev & test sets

It is critically important that worknotes such as Excel files or Tableau workbooks which include aggregate totals and charts which relate to detailed insights be prepared by an experienced machine learning specialist in conjunction with a data scientist in such a way that the analyses can be re-produced to ensure that the source data has not been tampered with (this is a security measure).

3. **(Re)Define:** This is where one (*II*) identifies a corresponding parameterization
 - 'translation' of the scientific question into statistical terms
 - Conduct a Bayesian group-think session for translation of substantive knowledge and quantities
 - output is a better defined framing of the problem domain

You should also consider at this stage of the process adjusting objective functions, targets, risk preferences & exploration vs. exploitation posture by assessing value of information (VOI).

4. **Feature Selection:** This is where one (*III*) specifies prior distributions for the unknown parameters
 - can come from data or knowledge

- establish several options towards a sensitivity analyses
- identification of 'non-informative' priors

5. **Predict & Test:** This is where one (*IV*) characterizes the linkage between the parameters and the observed data:

- design features amenable to Bayesian analysis
- interpret the problem domain as a causal graph of random variables
- specification of the likelihood, loss & cost functions

and (*V*) applies Bayes Theorem:

- $\text{posterior} \propto \text{prior} \times \text{likelihood}$
- turn the Bayesian 'handle' with bayesian updates

Andrew Ng specifically enumerates a process for this stage, as follows:

- Get data (train, dev, test, human panel sets)
- Design model (according to framing)
- Train model (using labelled training data)
- Test model (using held out dev set)

6. **Evaluation:** This is where one (*VI*) interprets the results and refines the scientific question

- examine features of the posterior
- review learned random variable parameters and Bayesian structure
- Assess training set performance against held out dev & test sets

It is critically important that this step be carried out by a trained machine learning specialist, and that numerous performance metrics be calculated and the scripts required to re-run the experiments be stored such that they can be re-run to check that the model has not been tampered with (as a security measure).

7. **Launch Approval:** Should only proceed if a human panel has been out-performed and a necessary, sufficient & complete Turing Test has been passed (if required).
8. **Execute:** If a model passes launch approval, it may be used to make a single coarse grained decision worth many millions of dollars, or many millions of small high frequency decisions each worth a few cents.