

Research Brief

Framing Requirements for Predictive Analytic Projects with Decision Modeling

August 2015

Written by: James Taylor

Key Takeaways

1. Organizations are struggling to create a scalable, sustainable operating model for decisions driven by predictive analytics.
2. In particular, firms need to a) ensure that predictive analytic projects address real business problems, b) engage a broad business audience in defining these problems, c) prioritize projects based on these definitions and d) ensure their analytic efforts deliver solutions to these problems.
3. Decision modeling, specifically decision modeling using a newly standardized notation, is an approach that explicitly defines the decision that will be improved by the predictive analytic before beginning the project.
4. By identifying the business problem in this way, organizations can ensure that predictive analytic efforts can be compared and prioritized while creating a shared understanding of the business problem and improving the odds of a successful implementation of the predictive analytics developed.

Business Problems

The effective use of predictive analytics is a topic that is exploding in importance across industries and in organizations of all sizes. As organizations adopt predictive analytics, and particularly as their predictive analytics projects move “out of the lab,” a consistent pattern of challenges can be observed:

- Organizations need a way for business stakeholders to frame the right questions to solve their business problems and set the “proper” requirements for analytic projects.
- Broad adoption of predictive analytics requires the engagement of people outside the data science/analytics team in thinking about how predictive analytics can solve business problems.
- Organizations need a way to pick the “right” analytics projects by prioritizing across a portfolio of projects.
- They also need a way to spot the “red flags” that show that an analytic project is likely to fail to deliver business value.
- Many are seeking a standard approach or formula to identifying the value of an analytic project and for communicating that value.

The root problem underlying all these challenges is one of framing a business problem so it can be effectively solved with predictive analytics. Without a standard, effective approach to framing the business problems involved, organizations will struggle to focus their predictive analytic investments effectively.

These challenges can be put in context using the most commonly used methodology framework for predictive analytics, CRISP-DM. Applying this model it becomes clear that these challenges are ones of *Business Understanding*. If a predictive analytic project does not have a clear understanding of the business problem then it is impossible to understand or prepare the right data, identify the model that will be required, and evaluate the ability of the model to meet the business need. In addition, deploying predictive analytic models is often a barrier to analytic success with “good” models failing to be deployed into good business solutions that meet the business need.

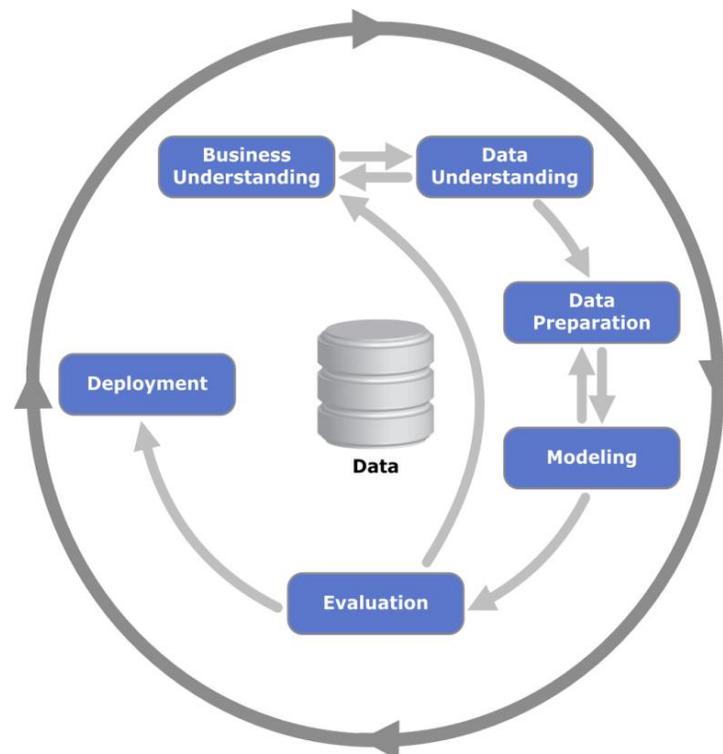


Figure 1. The CRISP-DM Lifecycle

Decision Modeling Addresses These Problems

An approach that is gaining significant traction and showing great promise in framing analytics projects is that of decision modeling. Decision modeling is used in the business understanding stage of a predictive analytic project. This technique, and a newly formalized standard notation, are the focus of this brief.

Introducing Decision Modeling

Decision modeling creates a formal, defined model of a decision-making approach using a standard notation. This notation, the Decision Model and Notation standard from the Object Management Group,¹ defines a new diagram type for specifying decision-making. A Decision Requirements Diagram specifies decisions and their component sub-decisions, the data that must be input to those decisions, and the knowledge required to make those decisions. This knowledge can be based on policy, regulations, or best practices, but it can also be analytic knowledge (especially predictive analytic knowledge) derived from data. These diagrams show the structure of the decision-making approach that is being followed, should be followed in the future, or is desirable as a best practice.

¹ <http://www.omg.org/spec/DMN/Current>

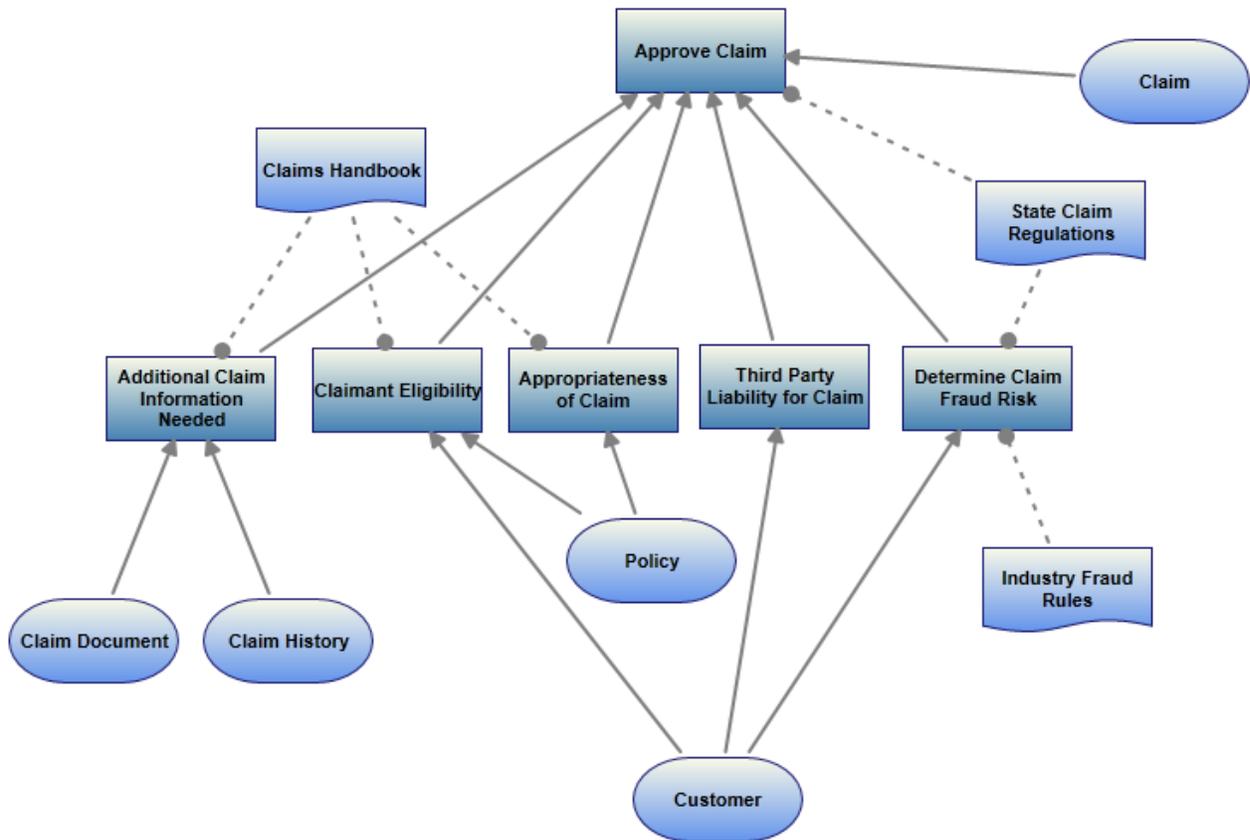


Figure 2. Decision Requirements Diagram Example from Insurance Industry

In the example above, for instance, the decision to approve a claim is shown as requiring information about the claim, access to State claims regulations and the claims handbook as well as decisions about the need for additional information, claimant eligibility, claim appropriateness, potential third party liabilities and fraud risk.

The standard also outlines how the decisions specified in the model can also be linked to business metrics or objectives, business processes and systems, as well as organizational units or roles. These links clarify the business context of the decision.

Predictive Analytic Requirements

Building a decision requirements model to specify business understanding at the very beginning of a predictive analytic project allows the creation of predictive analytic requirements that:

- **Describe a clear target for the project.** The decisions that the predictive analytic will influence are specified. The decision requirements diagram links these decisions to the ultimate business metrics or objectives that will be impacted by the analytic.
- **Identify the analytics to be developed.** Each piece of analytic knowledge can be described along with the information to be analyzed to produce it.
- **Is specific about which decisions are being influenced.** The decision requirements diagram shows which part of the decision making is influenced, exactly, by each predictive analytic model being developed and what other factors influence that decision-making.
- **Is specific about deployment.** The links for the decisions involved show which organizations will be involved, which business processes will be impacted and which systems will have to be altered.

A complete set of requirements for a predictive analytic project should include other project details such as executive sponsor, timeline, resources, planned analytic approach, etc. The decision requirements model allows the business problem being addressed by the project to be described more precisely, but it does not replace these other elements.

Implementing Decision Modeling

Predictive analytic teams that want to implement decision modeling as a requirements approach can follow the approach defined below. While the approach does not require a change in analytic technology, there are some technology considerations and early adopters have already identified some lessons learned (these will be expanded in a forthcoming brief).

Approach

1. Begin with performance measures

The first step for a predictive analytic project adopting decision modeling is to identify the performance measures, metrics, or business objectives for the project. Most organizations already identify the objectives of their predictive analytic project as well as some performance measures that will let them tell if their project was successful. It is critical that these are *business* metrics or objectives not analytic ones. For instance business metrics such as reducing the number of claims in litigation and improving the loss ration rather than analytic metrics such as delivering a model with a certain error rate or lift.

2. Identify decisions that matter

Having identified the measures of success, analytic teams can then identify the various decisions the organization makes that impact this measure or set of measures. A predictive analytic model cannot impact a measure directly, it can only improve a decision such that improved decision making will impact that measure. A predictive analytic model, for instance, cannot improve an organization's loss ratio. It can improve the decision "should this claim be approved for payment?" by identifying claims likely to be fraudulent and by doing so improve the loss ratio.

Predictive analytic models are particularly effective at improving operational decisions – decisions about a single customer or single transaction. The nature of predictive analytic models is that they predict something about someone or something. A claim fraud model, for instance, predicts how likely THIS claim is to be fraudulent. Finding decisions about the same topic – a claim – helps ensure that the decision can be improved using the predictive analytic. These decisions also repeat – they must be made over and over – making it possible to define and manage a model of the decision making to be used each time.

3. Clearly define decision(s) to be improved

Once it is clear which decision or decisions will be targeted for improvement with the predictive analytic being developed, the analytic team needs to ensure that there is an agreed definition of the decision. While a name and description can be effective in defining the decision, it is often the case that there is disagreement as to exactly what the decision involves. This disagreement may be explicit or implicit: There may be explicit disagreement between people or organizations as to how the decision should be made or inconsistency between decision-makers today. More likely there is implicit disagreement where different perspectives or assumptions mean that people do not really agree on how the decision should be made even though they think they do.

To ensure that there is agreement teams should develop a question and a set of allowed answers for the decisions being targeted. Questions work well for describing decisions as making a decision involves making a choice, selecting one of several possible outcomes. This can be clearly specified as a question. The question should be very specific with respect to the subject of the decision, its timing and its scope. "How" and "We" should be avoided and a question such as "Which retention offer should the company extend to this customer when they call to cancel their service?" is much better than "How can we retain this customer?" The possible or allowed answers to the question should also be specified. Allowed answers range from Yes/No to Accept/Reject/Refer to "Any currently active marketing offer defined in the database." Making sure everyone is clear what the end result of the decision must be – what

the possible answers are – helps ensure that any predictive analytic developed will actually help by improving the ability of the organization to discriminate between these possible answers.

For instance an “Approve Claim?” decision is not, in fact, a Yes/No decision in most organizations. More likely it can be defined as follows:

In what way should this claim be processed once complete details have been received?

- Auto-approve
- Fast-track
- Regular process
- Reject
- Fraud investigation unit

This makes it clear that there are degrees of approval and that it is important to identify which approach is most suitable for a given claim.

4. Specify business context

A clearly defined decision can and should be put into a strong business context. This business context should include:

- **Other measures and metrics impacted.** Many decisions involve making tradeoffs between several metrics or measures, while others have a wide ranging impact across an organization’s metrics. Understanding this impact ensures that improvements are true improvements not local optimizations.
 - For instance the Number of Claims Opened/Closed per month and Average Claim Payment metrics might also be impacted by the Approve Claim decision for instance.
- **Processes and systems impacted.** Identifying the business processes and systems in which the decision-making is required will establish where and how the decision must be made and so inform the deployment of any predictive analytic model developed. For instance the approve claim decision might be part of the Settle Claim business process and be embedded in a legacy claims system.

- **Organizational units and roles involved.** Understanding which organizational units or roles must approve the decision-making approach (the owners of the decision) as well as which have actually make it ensures that the analytic team engages with the right people. Predictive analytics will have to be used, understood and approved by people outside the-analytic team. Identifying these people early will save confusion and re-work later.
 - For instance, while Claims Adjustors make the decision to approve claims, the approach used is owned by the Claims department as a whole. Actuarial and Fraud groups are clearly going to be impacted by the decision also.

Developing a business context for the decision will help ensure that deployment issues are identified early, making it more likely that the predictive analytic will be deployed and used. It will also help ensure that a balanced perspective of the decision to be improved is available when the predictive analytic is developed.

5. Define data requirements

With a sense of the business context the team can now start developing the decision requirements diagram itself. The first step in developing the decision requirements diagram is to identify the data required to make the decision. This is referred to as Input Data as it is the data or information that must be input to the decision-making context from outside if the decision is to be made effectively. The data required by a decision is generally going to be:

- Transactional data – the “thing” that is the subject of the decision such as data about the customer.
- Case data – data related to the transactional data such as data about products the customer owns.
- Reference data – data that is used as part of the decision-making approach.

In each case the business entity is identified as a piece of Input Data and shown as an oval. These are linked to the decision (shown as a rectangle) using a solid arrow known as an information requirement. Decisions require a piece of input data if some or all of the information represented by that input data is used in the decision-making.

Decision requirements models can be developed for all or part of an identified business decision. Part of the claim approval decision described above relates to fraud. The claim’s fraud risk must be assessed by answering the question:

Does this claim seem suspicious and, if so, how suspicious does it seem?

- Not suspicious
- Somewhat suspicious
- Very suspicious

An initial analysis of this decision might identify that someone making this decision requires information about the Customer, the Claim, Claim Documents submitted and the Claim History.

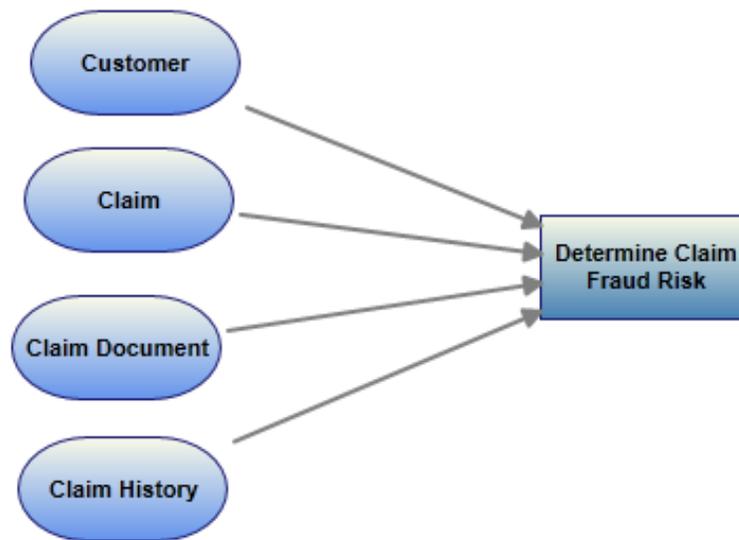


Figure 3. Defining Data Requirements Example

6. Define knowledge requirements

The next step is to identify the knowledge required to make the decision – how to use this input data to determine the correct answer. This knowledge can be of various types:

- External regulations that constrain how the decision is made.
- Internal policies that describe the general decision-making approach.
- Expertise and best practices that define how the decision has been made to date.
- Analytic insight that improves the accuracy of the decision or might do so in the future.

To identify this knowledge the team can consider what tells someone making the decision what they must do, what they should do, what they can do or what they will probably do. Each such source of knowledge is defined as a Knowledge Source and shown as a document shape. The decision or decisions influenced by each piece of knowledge are linked to them using a dashed round-headed link known as an authority requirement. A decision requires a knowledge source as an authority if the knowledge source describes, enables, improves or constrains the decision-making.

The predictive analytic model being developed will be shown as a knowledge source in this way. If it is not clear what analytic should be developed, or if the team is considering various approaches, it can be useful to ask the question “What would help me make this decision more accurately?” or to get business owners to complete the phrase “If only we knew xyz we could make a more profitable decision”.

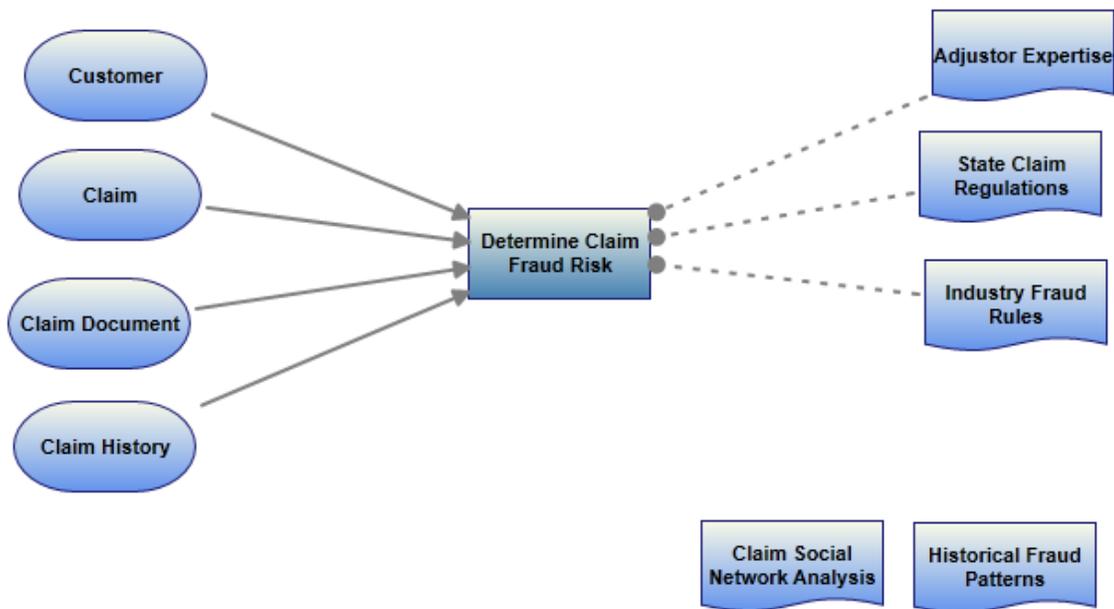


Figure 4. Defining Knowledge Requirements Example

For instance the current state might be that fraud is detected by applying the adjustor’s expertise, state claim regulations and industry fraud rules. The analytics project might identify that social network analysis and historical fraud pattern identification would be helpful. The question is how, exactly, can they help make this decision?

7. Define decision requirements

The next step is to identify additional decision-making requirements. The reality of most significant business decisions is that there are other decisions that must be made first – sub-decisions or precursor decisions. It is not possible to make a significant decision about a customer, say, before several other decisions have been made about that customer.

For example interviewing claims experts might reveal that claims fraud can be divided up into suspicions about the claim itself and suspicions about the claimant. Some providers might also be known as suspicious. Deciding if the claim should be considered OK, somewhat suspicious, or very suspicious will require each of these other decisions to be made first.

If a decision requires another decision in this way it depends on that decision being made first – that other question being answered. When a decision requires another decision it is shown using the information requirement link as it is the information that results from making the decision that is required.

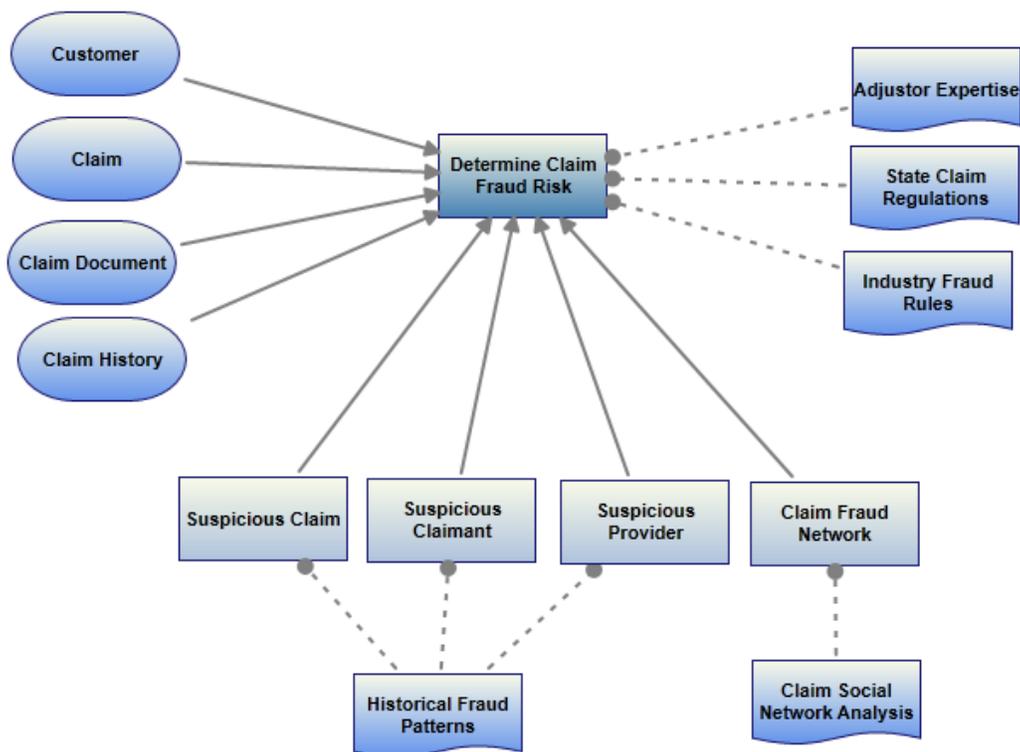


Figure 5. Defining Decision Requirements Example

This decomposition immediately makes it clear which part of the fraud decision will be impacted by which kind of analysis. In addition, because the experts have already identified that each of these four questions is a Yes/No question, the way in which analytics might help is both known and constrained.

In a real project claims adjustors were used to making Yes/No decisions about the fraud level in a claim. The analytics team built a more effective model for detecting fraud but one that placed claims into quintiles by fraud risk. Operationalizing this analytic was unexpectedly difficult as the analytic team had not realized that they were imposing a change in decision-making approach. Once a decision requirements modeling exercise revealed this (and some other related challenges) the model was successfully rolled out.

To find decision requirements the team can consider:

- Where does the decision get information from?
 - Some of the information that is required is not available in the environment before decision-making starts but must be determined by making another decision first. For instance there is no flag on the Provider record that the Provider should be considered suspicious. This must be decided based on other data.
- Is the decision “atomic”?
 - If part of the decision can be made without making the rest then it may be possible to define the logical components of the decision. If someone makes the decision currently then their description of their decision-making approach will likely, naturally, describe it in steps or pieces that can be identified as sub-decisions.
- Do circumstances change the decision making?
 - Some decisions are made differently, or not made at all, if specific circumstances are true and will therefore depend on decisions that determine those circumstances.

8. Iterate and refine

Building a decision requirements diagram is a highly iterative process. As each new decision is identified it should be described with a question and allowed answers, then assessed to see what input data and knowledge it requires. These may already be in the diagram, but often represent new, more specific information and authority requirements.

As the diagram is extended it may become clear that specific pieces of input data or specific pieces of knowledge are not required by the decision to which they are linked, but by one of its sub-decisions. This allows for more specificity as to the role of that input data or knowledge source. For instance it may become clear that Industry Fraud Rules relate only to identifying suspicious providers or that part of identifying a claimant as suspicious is a review of the kind of policy they have purchased.

Input data can also be assessed to see if it is really used in its natural state or if it must be processed first – if there is a decision that must be made with that data first.

Decision requirements diagrams can be reviewed with business, IT and analytic teams to ensure that everything is understood correctly. Eventually the team will have a model that accurately reflects the current or intended decision-making. One that shows clearly where and how the proposed predictive analytic(s) will influence that decision-making.

The role of each piece of analytic knowledge is clearly shown in the diagram, making it clear how the analytic will improve the decision-making. Similarly the role of other kinds of expertise is clear. For instance, the example model below identifies that adjustors will still have a say in how the final fraud determination is made but that they will not have a say in how each element of the decision is made – these will be driven by other kinds of knowledge.

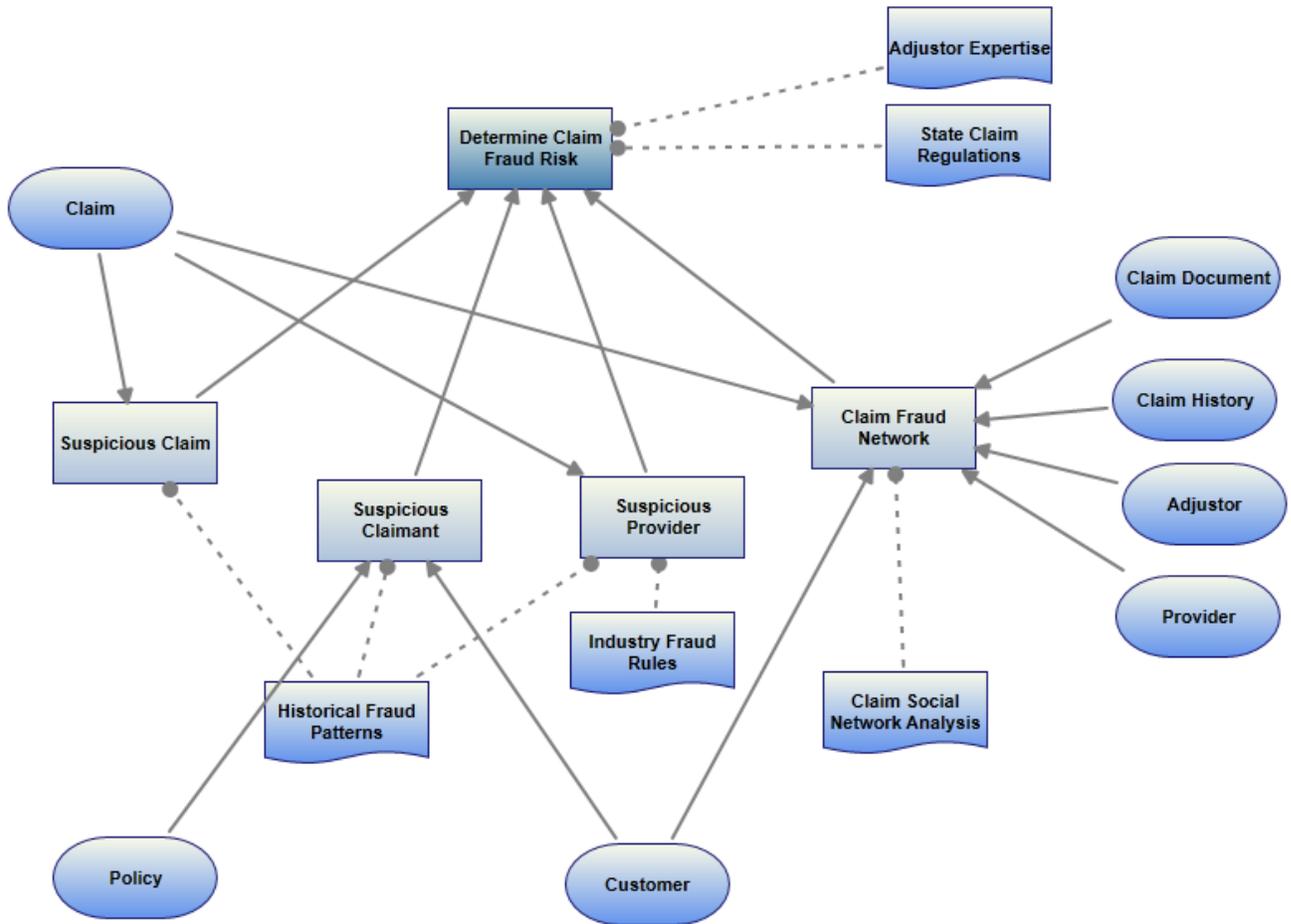


Figure 6. Decision Requirements Document

The notation can also be used to develop an initial set of analytic data requirements, showing which kinds of data are expected to contribute to which analytic knowledge. The input data from which the analytic knowledge is derived are linked to that knowledge source using authority requirements links as shown below.

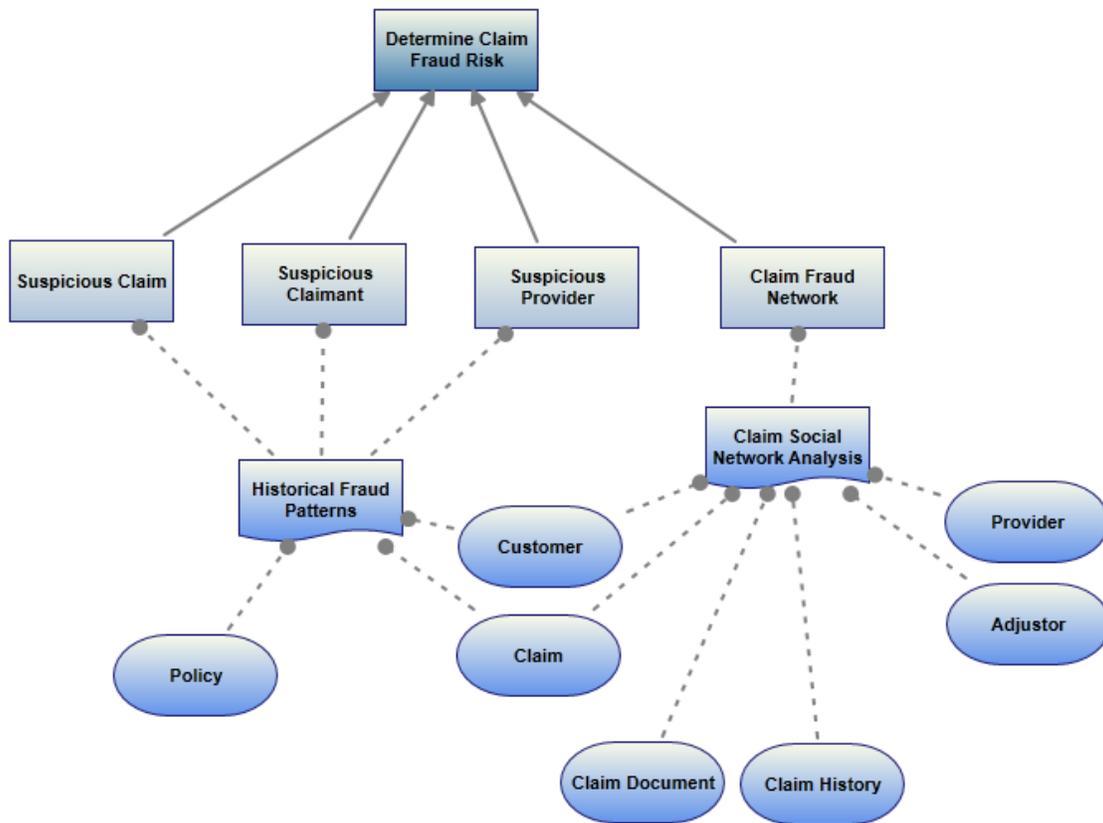


Figure 7. Analytic Data Requirements Example

Technology Considerations

Developing decision requirements models for analytic projects does not require analytic technology per se. The models use a simple notation and can be developed using standard drawing and documentation tools. It is worth considering a proper modeling platform, however, to maintain the models developed in a shared repository. Such a tool can support collaborative model development by potentially distributed teams, allow decisions as well as

analytics and input data to be reused on subsequent projects, and facilitates a longer term view of decision-making and its systemic improvement.

Lessons Learned

While the use of decision modeling in analytic projects is still in its infancy, some clear lessons have already been learned by early adopters:

- It is essential that business metrics and objectives are used, not simply analytic ones. Unless the results of analytic projects are clearly expressed in business terms, analytics will remain a technical, siloed activity.
- There should be no “white space” between analytic success and business success. That is to say that a successful analytic project should be a successful business project. Accurate predictive models that are not being used or are hard to deploy are business failures and must be considered analytic failures also. Decision modeling helps reduce the white space by making it clear how the analytic can be used.
- Suitable decisions need to be selected. Repeatable decisions that the organization wishes to make consistent, to some degree at least, are ideal. Only if there is a willingness to standardize some elements of the decision-making can the potential for analytics be seen. Decision modeling makes the relative roles of expertise, regulation and analytics explicit allowing a transparent discussion of these issues.

Example Results

As more organizations use decision modeling on predictive analytic projects more positive results are becoming apparent. While a Leading Practices Case Study is planned for later in 2015, some specific example results include:

- A leading insurance provider developed a decision model for its healthcare claims adjudication decision. The model showed where and how a fraud analytic could be used materially improving the roll out of this analytic and increasing usage.
- A life insurance provider used a decision requirements model to identify that a planned analytic to predict medical risk would not be usable, allowing the analytic project to be redirected before it began.

- A manufacturing company identified that while predictions of quality risk could not be used to change task assignments for organizational reasons, they could be used to focus the daily activities of supervisors and the quality team.

Conclusion

A decision requirements model acts as an effective statement of business understanding. By clearly defining the decision-making to be analytically improved and the role of that decision making in the broader business context, a decision requirements model:

- Ensures that business stakeholders and analytic professionals alike know the right question to answer analytically.
- Allows business analysts and business stakeholders to identify analytical requirements in their projects before engaging data science and analytic teams.
- Allows analytic projects to be compared with each other so that those most likely to have a significant business impact can be prioritized.
- Shows red flags such as poorly understood decision-making, complex regulatory environments, deployment challenges and organizational risks before analytic development begins.
- Standardizes the definition of analytic value in terms of the decisions to be improved and their impact on known business metrics and KPIs rather than in terms of analytic model accuracy.

About the Author

James Taylor is the CEO of Decision Management Solutions. James is the leading expert in Decision Management Systems— systems that are active participants in improving business results. Decision Management Systems apply business rules, predictive analytics and optimization technologies to address the toughest issues facing businesses today, changing the way organizations are doing business. James is passionate helping companies develop agile, analytic and adaptive processes and systems. James has over 20 years developing software and solutions for clients and has led Decision Management efforts for leading companies in insurance, banking, health management and telecommunications.

In addition to strategy and implementation consulting, James is an experienced and highly rated keynote and speaker at conferences in the US and around the world. James also regularly runs webinars for clients, and as educational outreach for Decision Management Solutions. James wrote “Decision Management Systems: A Practical Guide to Business Rules and Predictive Analytics” (IBM Press 2012). He recently published “The Microguide to Process and Decision Modeling” with Tom Debevoise and previously wrote “Smart (Enough) Systems” (Prentice Hall, 2007) with Neil Raden. He has contributed chapters to several books on business rules and business analytics.

References

- ▶ Debevoise, Tom and Taylor, James (2014). *The MicroGuide to Process and Decision Modeling in BPMN/DMN: Building More Effective Processes by Integrating Process Modeling with Decision Modeling*.
- ▶ Decision Management Solutions (2015), *Decision Modeling with DMN*.
- ▶ Decision Management Solutions (2015), *Framing Analytic Requirements*.
- ▶ Object Management Group. *Decision Model and Notation (DMN) Specification 1.0*
Current version at <http://www.omg.org/spec/DMN/Current>
- ▶ Taylor, James (2012). *Decision Management Systems – A Practical Guide to Using Business Rules and Predictive Analytics*. IBM Press.