

1. Comprehensive Introduction to Intelligent Software Agents

This is a dangerous time for professional accountants. Technology is changing very rapidly¹. For example, terms like “Big Data” and “Deep Learning” and “Artificial Intelligence” are being thrown about by software vendors, some of which are little more than snake oil salesman; with these new technologies comes a significant amount of hype. On the other hand, many of these software vendors have very useful products what will transform financial reporting as it is practiced today. Knowledge is very useful when one is trying to differentiate the hype from real possibilities.

Setting the right expectations is another big challenge. This example will help you understand what I mean. There is a lot of good and bad information relating to driverless cars. Tesla has functionality in its cars that is similar to driverless car-type functionality. Imagine trying to set a goal for what you want such a car to achieve:

- **Goal A:** Create a car that can travel anywhere in the United States without a driver.
- **Goal B:** Create a car that can assist the driver to drive anywhere in the United States.
- **Goal C:** Create a car that can assist a driver when driving on an Interstate Highway.

What is the relative cost of achieving goals A, B, and C above? What is the relative probability of even being able to achieve goals A, B, and C without significant, expensive research and development?

The point is that setting the right expectations helps one understand what is actually practical and useful. Thinking that technology will have no impact on how you perform your work will be a complete disaster for your career as an accountant and/or for your business.

A core part of expert systems and other knowledge based systems is the intelligent software agent. This document helps professional accountants sort through all the information and misinformation that they are hearing related to XBRL-based structured digital financial reporting and helps them understand what an intelligent software agent is, basics on how they work, and examples of how they might be used.

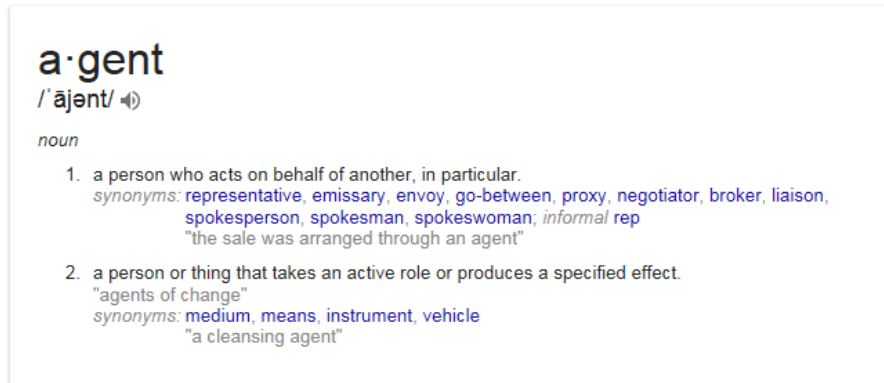
1.1. Deconstructing the Notion of an Intelligent Software Agent

To understand intelligent software agents one first needs to understand the notion of an agent. This section is dedicated to setting your perspective as to the notion of an agent. The section provides specific definitions, deconstructing the pieces so that we can subsequently put the pieces back together.

¹ AICPA News Update: Technologies are poised to reshape the accounting landscape, <http://xbrl.squarespace.com/journal/2017/7/7/aicpa-news-update-technologies-are-poised-to-reshape-the-acc.html>

1.1.1. Definition of an agent

Google defines agent² as “a person who acts on behalf of another” and “a person or thing that takes an active role or produces a specified effect”.



An agent performs specific tasks on behalf of another.

Important to the definition of agent is the law of agency³. Some principle, or the person making use of the agent, authorizes the agent to work on their behalf to perform an agreed upon specific task. The important key points are the *authority* to act and an *agreement* on the specific task or tasks that the agent is supposed to perform.

Business professionals generally understand the notion of an agent because they get introduced to this notion in their Business Law 101 class in college. This class introduces them to the notion of human agents. But what if the agent is not a human?

1.1.2. Software agents

In the case of computer science, an agent is a software program. Wikipedia defines software agent⁴ (paraphrasing) as “a computer program that acts for a user or other program in a relationship of agency to perform some action”.

Software agents are sometimes also referred to as “bots” from the term robot.

1.1.3. Intelligent agents

An intelligent agent⁵ is an abstract notion that links the *real world* agent and the notion of agency with an *implementation* of that functionality within software. An intelligent agent is the abstract functionality of a system similar to a computer program; it is not the computer software program itself.

Intelligent agent, as we are using it, is an idea related to artificial intelligence. An intelligent agent is an autonomous entity which observes its environment through sensors and acts upon that environment using actuators in the pursuit of some goal. Intelligent is a software design philosophy.

² Google search, *Agent definition*, retrieved August 14, 2016, <https://www.google.com/search?q=agent+definition>

³ Wikipedia, *Law of agency*, retrieved August 14, 2016, https://en.wikipedia.org/wiki/Law_of_agency

⁴ Wikipedia, *Software agent*, retrieved August 14, 2016, https://en.wikipedia.org/wiki/Software_agent

⁵ Wikipedia, *Intelligent agent*, retrieved August 14, 2016, https://en.wikipedia.org/wiki/Intelligent_agent

1.1.4. Intelligent software agents

The main difference between a software agent and an ordinary program is that a software agent is autonomous; that is, it must operate without direct intervention of humans or others.

Intelligent software agents are engineered using specific principles⁶.

1.1.5. Notion of an intelligent software agent which we will use

An agent is an entity capable of sensing the state of its environment and acting upon it based on a set of specified rules.

The consultancy firm McKinsey predicts that intelligent software agent use will increase in business⁷. Use of intelligent software is expected to be the norm in law firms⁸ by the year 2020.

1.2. Understanding what Intelligent Software Agents Do

This section helps you understand what intelligent software agents do and how they do it.

1.2.1. Basic example of an intelligent software agent

A basic example of an intelligent software agent is a basic thermostat^{9,10}.




Fig. 1: Thermostat agent

Software reads the current room temperature using a sensor and uses a set of rules to turn the heat on or turn the heat off to impact the future room temperature in the room or environment in which the sensor is located.

1.2.2. Basic model of an intelligent software agent

If you look at the basic intelligent software agent above and then you look at lots of other tasks that such intelligent software agents do you see patterns. Intelligent

⁶ Frances M.T. Brazier, Catholijn M. Jonker, Jan Treur
Vrije, *Principles of Component-Based Design of Intelligent Agents*,
<http://www.few.vu.nl/~wai/Papers/DKE02.princ.pdf>

⁷ McKinsey, *Artificial Intelligence Meets the C-Suite*, [http://www.mckinsey.com/business-](http://www.mckinsey.com/business-functions/strategy-and-corporate-finance/our-insights/artificial-intelligence-meets-the-c-suite)
[functions/strategy-and-corporate-finance/our-insights/artificial-intelligence-meets-the-c-suite](http://www.mckinsey.com/business-functions/strategy-and-corporate-finance/our-insights/artificial-intelligence-meets-the-c-suite)

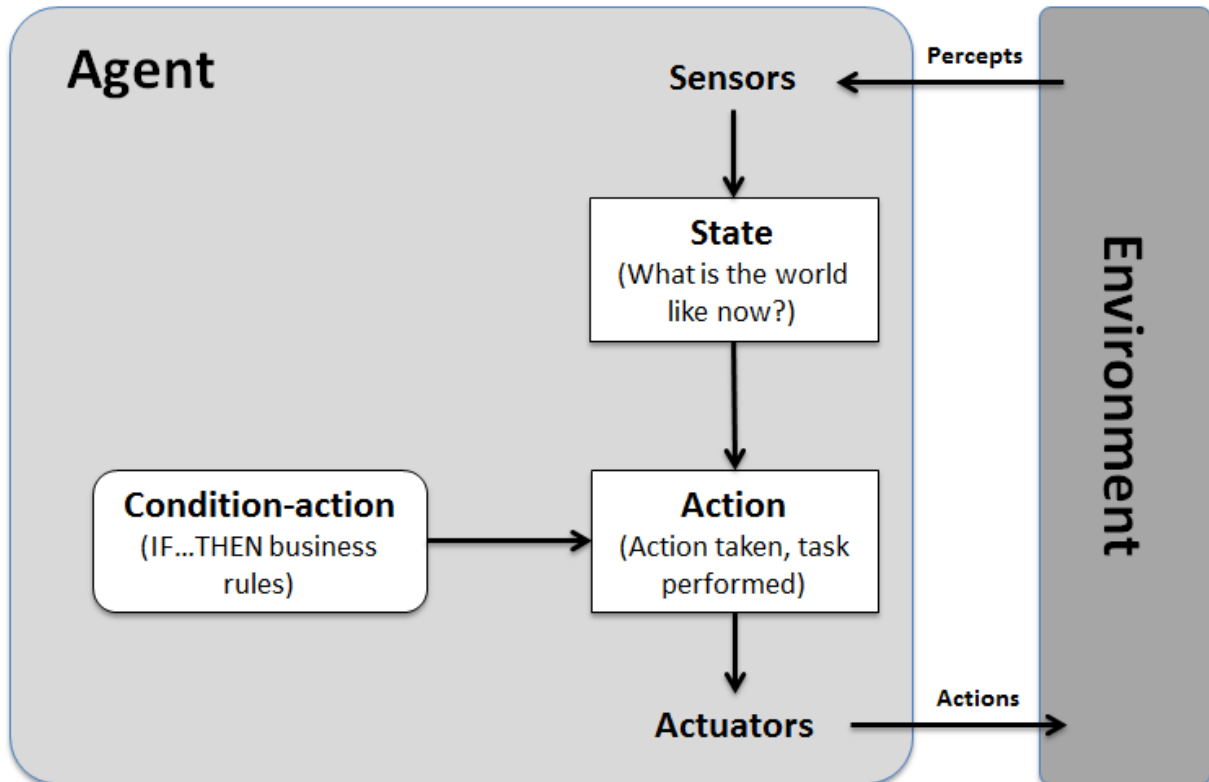
⁸ Managing Partner, *Artificial intelligence to be 'the norm' in law firms by 2020*,
[http://www.managingpartner.com/news/business-strategy/artificial-intelligence-be-%E2%80%98-](http://www.managingpartner.com/news/business-strategy/artificial-intelligence-be-%E2%80%98-norm%E2%80%99-law-firms-2020)
[norm%E2%80%99-law-firms-2020](http://www.managingpartner.com/news/business-strategy/artificial-intelligence-be-%E2%80%99-law-firms-2020)

⁹ Consortium of Cognitive Science Instruction, *Introduction to Intelligent Agents*, Retrieved July 24, 2016;
http://www.mind.ilstu.edu/curriculum/ants_nasa/intelligent_agents.php

¹⁰ Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach (Third edition)*, page 31,
<http://aima.cs.berkeley.edu/>

software agents leverage these patterns. Intelligent software agents are computer code written in a specific way. An **agent** is an entity capable of **sensing** the **state** of its **environment** and **acting** upon it based on a **set of specified rules**. Remember that an agent performs specific tasks on behalf of another in order to achieve some agreed upon goal established by the principal that employs the agent and the agent which will act in behalf of the principle.

In the case of software, an agent is a software program. Consider that definition of an agent and look at the graphic below to get an idea of how intelligent agent software works:



An intelligent agent is software that assists people and acts on their behalf. Intelligent agents work by allowing people to:

- delegate work that they could have done to the agent software,
- perform repetitive tasks,
- remember things you forgot,
- intelligently find, filter and summarize complex information,
- customize information to your preferences,
- learn from you and even make recommendations to you.

1.2.3. Business rules drive intelligent software agents and expert systems automating work

An example of one major shift is provided by what professional accountants call the "disclosure checklist". Accountants creating financial reports often use accounting

and reporting checklists or "disclosure checklists" as memory joggers to help them get the reports right¹¹. These memory joggers were created to be read by humans and can be a couple hundred pages. What if a financial report was structured, such as an XBRL-based public company financial report that must be submitted to the U.S. Securities and Exchange Commission. What if these human-readable memory joggers could be made machine-readable. And what if an intelligent software agent could be created to automate the manual task of checking a financial report to make sure that report was mechanically correct.

Note the statement "mechanically correct". This is a very important distinction. No computer program will ever have the judgement of a professional accountant. See the section *Setting the right expectations* later in this document. But computer programs can perform work if the financial report is structured and the necessary business rules are made machine-readable. How much of a disclosure checklist can be automated? That percentage is as-of-yet to be determined. Perhaps 20% can be automated or even 80% will be automated. Maybe even a higher percentage. The probability 0% of a disclosure checklist can be automated is extremely low.

Humans augmented by machine capabilities, much like an electronic calculator enabling a human to do math quicker, will empower knowledge workers who know how to leverage the use of those machines.

1.2.4. Important terminology related to intelligent software agents

The following is a summary of important terminology¹² that is used to discuss intelligent software agents.

- **Agent:** An **agent** is an entity capable of **sensing** the **state** of its **environment** and **acting** upon it based on a **set of specified rules**. Remember that an agent performs specific tasks on behalf of another in order to achieve some agreed upon goal established by the principal that employs the agent and the agent which will act in behalf of the principle.
- **Environment:** environment in which the agent operates; description of the state of affairs that change over time as real world situations do.
- **Sensing capabilities:** capability of the agent to understand its environment; determines the sort of data the agent is capable of receiving as input.
- **Percept:** A percept refers to the agent's perceptual inputs at any given moment.
- **Percept sequence:** The percept sequence represents the complete sequence of percepts the agent has sensed or perceived during his lifetime.
- **State:** current conditions of the environment. Or, a projected conditions of the environment based on actions taken.
- **Actions:** change in the environment brought about by the agent, requiring the agent to update its model of the world, which in turn may cause the agent to change its immediate intention.

¹¹ *Automating Accounting and Reporting Checklists*, <http://xbri.squarespace.com/journal/2016/5/5/automating-accounting-and-reporting-checklists.html>

¹² Consortium of Cognitive Science Instruction, *Introduction to Intelligent Agents, Taxonomy of agents*, Retrieved July 24, 2016; http://www.mind.ilstu.edu/curriculum/ants_nasa/intelligent_agents.php

- **Condition-action rules:** A formal and implementable expression of some business user requirement. Includes definitions of terms, structural assertions, action assertions, and derivations.
- **Desires:** overall policies or goals of the agent.
- **Action selection architecture:** an agent decides what to do next by consulting both its internal state, the state of the world, and its current goal; then the agent uses decision making procedures to select an action.

1.2.5. Agents Working Together: Multi-agent Systems

Note that one agent must be able to communicate with other agents and exchange information between agents. Below is terminology related to the communication between agents¹³.

- **Multi-agent system:** When an agent coexists in an environment with other agents, perhaps collaborating or competing with them, the system is considered a multi-agent system.
- **Coalition:** A coalition is any subset of agents in the environment.
- **Strategy:** A strategy is a function that receives the current state of the environment and outputs the action to be executed by a coalition.
- **Blackboard:** A blackboard structure is a communication form that consists of a shared resource divided into different areas of the environment where agents can read or write any significant information for their actions.
- **Coordination:** Coordination is essential in multi-agent system because it provides coherency to the system behavior and contributes to achieving team or coalition goals.
- **Cooperation:** Cooperation is necessary as a result of complementary skills and the interdependency present among agent actions and the inevitability of satisfying some success criteria.
- **Competition:** Another possible model is that in which the agents are self-motivated or self-interested agents because each agent has its own goals and might enter into competition with the other agents in the system to achieve these goals. In this sense, competition might refer to accomplishing or distributing certain tasks.
- **Negotiation:** Negotiation might be seen as the process of identifying interactions based on communication and reasoning regarding the state and intentions of other agents.

When one intelligent software agent communicates with another intelligent software agent some technical syntax and semantics must be employed¹⁴. The XBRL global standard can be used for such communication and collaboration.

¹³ C# - Applying AI to a Multi-Agent 'Mini-Basketball' Game, Retrieved July 24, 2016; <https://msdn.microsoft.com/en-us/magazine/mt736456.aspx>

¹⁴ Hanh Tran & Thaovy Tran, *Intelligent Agent, Agent Communication Languages*, http://groups.engin.umd.umich.edu/CIS/course.des/cis479/projects/agent/Intelligent_agent.html#dep6

1.3. *Building your Understanding of Intelligent Software Agents*

We will build your understanding of intelligent software agents piece-by-piece in this section.

1.3.1. *Artificial intelligence*

Artificial intelligence is a branch of computer science. There are many good descriptions of artificial intelligence¹⁵. Here is one good definition:

Artificial intelligence is the automation of activities that we associate with human thinking and activities such as decision making, problem solving, learning and so on.

Those trying to make artificial intelligence work over the past 40 or so years have had limited success. But that is changing. Both under estimating or over estimating the capabilities the computer software will be able to achieve can have catastrophic consequences.

There are two types of artificial intelligence, specialized and generalized:

- **Specialized:** An example of specialized artificial intelligence is programming a computer to play chess. The software performs one specific task. Specialized artificial intelligence is fairly easy to achieve.
- **Generalized:** An example of generalized artificial intelligence is the sort of science-fiction stuff you see in the movies. Why do you only see this in the movies? Because generalized artificial intelligence is extremely hard to make work.

1.3.2. *Expert systems*

Expert systems¹⁶ is a branch of artificial intelligence. The following is a definition of an expert system:

Expert systems are computer programs that are built to mimic human behavior and knowledge. The computer program performs tasks that would otherwise be performed by a human expert. A model of the expertise of a domain of knowledge of the best practitioners or experts is put into machine-readable form and the expert system reaches conclusions or takes actions based on that information.

1.3.3. *Business rules*

Key to employing artificial intelligence and therefore making an expert system or intelligent software agent¹⁷ work is business rules¹⁸ of the domain being put into machine-readable form.

¹⁵ AlanTuring.net, *What is Artificial Intelligence?*,

http://www.alanturing.net/turing_archive/pages/reference%20articles/What%20is%20AI.html

¹⁶ *Understanding the Components of an Expert System*,

<http://xbri.squarespace.com/journal/2016/5/24/understanding-the-components-of-an-expert-system.html>

¹⁷ Wikipedia, *Intelligent Agent*, retrieved July 24, 2016; https://en.wikipedia.org/wiki/Intelligent_agent

¹⁸ *Comprehensive Introduction to Business Rules*,

http://xbri.azurewebsites.net/2017/IntelligentDigitalFinancialReporting/Part01_Chapter02.4_ComprehensiveIntroductionToBusinessRules.pdf

1.3.4. Important distinctions between intelligent software agents

There are two important distinctions that you should keep in the back of your mind when we talk about intelligent software agents:

- **Rational agent:** A rational agent is one that acts so as to achieve the best outcome or, when there's uncertainty, the best expected outcome. Rationality as used here refers to following the rules of logical reasoning, making correct inferences, and selecting the appropriate action that will lead to achieving the desired goal.
- **Autonomous agent:** An autonomous agent is a system situated within and a part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to effect what it senses in the future.

Every intelligent software agent that we are interested in is generally both a rational and autonomous agent.

1.3.5. Categories of intelligent software agent functionality

The functionality of an intelligent software agent can be classified into groups, or categories:

- **Reactive agent:** A reactive agent is capable of maintaining an ongoing interaction with the environment and responding in a timely fashion to changes that occur in it.
- **Pro-active agent:** A pro-active agent is capable of taking the initiative; not driven solely by events, but capable of generating goals and acting rationally to achieve them.
- **Deliberative agent:** A deliberative agent symbolically represents knowledge and makes use of mental notions such as beliefs, intentions, desires, choices and so on. This is implemented using a belief-desire-intension model.
- **Hybrid agent:** A hybrid agent is one that mixes some of all the different architectures.

Every intelligent software agent falls into one of those four categories.

1.3.6. Categories of intelligent software agent sophistication

Intelligent software agents can be grouped as to the level of sophistication offered by the agent:

- **Generic agent:** An agent is anything that perceives an environment through sensors and acts or reacts upon the environment through effectors.
- **Simple reflex agent:** A simple reflex agent looks up what it should do from a list of rules in response to its perception to the environment.
- **Model-based reflex agent:** A model-based reflex agent is the same thing as a simple reflex agent except that a model-based reflex agent has a model of how the environment evolves.
- **Goal-based agent:** A goal-based agent has a goal or set of goals that it actively pursues in accordance with an agenda (so this type of agent is proactive, not just reactive). A goal based agent has a representation of the

current state of the environment and how that environment works. The agent pursues policies or goals that may not be immediately attainable. And so, goal based agents do not live merely in the moment. These agents consider different scenarios before acting on their environments, to see which action will probably attain a goal. This consideration of different scenarios is called search and planning.

- **Utility-based agent:** A utility-based agent is a more sophisticated type of goal-based agent that also rates each possible scenario to see how well it achieves certain criteria with regard to production of the good outcome, therefore it is more adaptive. A utility measure is applied to the different possible actions that can be performed in the environment. The utility-based agent will rate each scenario to see how well it achieves certain criteria with regard to the production of a good outcome. Things like the probability of success, the resources needed to execute the scenario, the importance of the goal to be achieved, the time it will take, might all be factored in to the utility function calculations.

Every intelligent software agent falls into one of those groups in terms of sophistication of functionality offered.

1.3.7. Learning versus non-learning intelligent software agents

Every intelligent software agent can be either a learning or non-learning agent:

- **Learning agent:** A learning agent is one that requires some training to perform well, adapts its current behavior based on previous experiences and evolves over time.
- **Non-learning agent:** A non-learning agent is one that doesn't evolve or relate to past experiences and is hard coded and independent of its programming.

A learning agent is significantly more sophisticated in terms of functionality, but the tradeoff is that they are also significantly more complex to create.

1.4. Intelligent software agents assisting humans

Artificial intelligence is the automation of activities that we associate with human thinking and activities such as decision making, problem solving, learning and so on¹⁹.

Expert systems is a branch of artificial intelligence. An intelligent agent is software that assists people and acts on their behalf. Intelligent agents work by allowing people to:

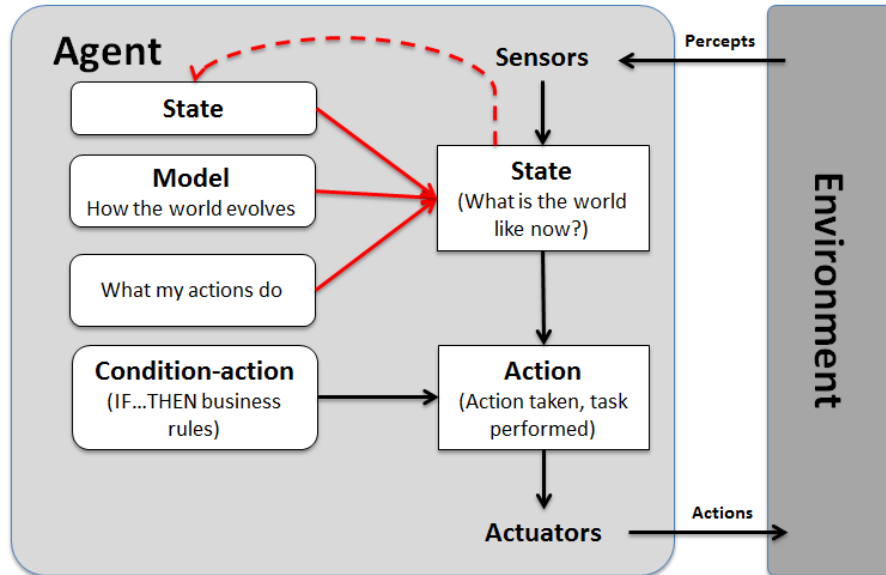
- delegate work that they could have done to the agent software,
- perform repetitive tasks,
- remember things you forgot,
- intelligently find, filter and summarize complex information,
- customize information to your preferences,

¹⁹ *Introduction to Artificial Intelligence Terminology*, <http://xbrl.squarespace.com/journal/2016/7/21/introduction-to-artificial-intelligence-terminology.html>

- learn from you and even make recommendations to you.

An **agent** is an entity capable of **sensing** the **state** of its **environment** and **acting** upon it based on a set of specified **rules**. An agent performs specific tasks on behalf of another. In the case of software, an agent is a software program.

The main difference between a software agent and an ordinary program is that a software agent is autonomous; that is, it must operate without direct intervention of humans or others. There are many different types of intelligent software agents²⁰.



Intelligent agents can perform sophisticated work. A rational agent is one that acts so as to achieve the best outcome or, when there's uncertainty, the best expected outcome. Rationality as used here refers to following the rules of logical reasoning, making correct inferences, and selecting the appropriate action that will lead to achieving the desired goal.

Machine-readable business rules are key to creating intelligent software agents that provide the functionality within an expert system.

1.4.1. Getting the thick layer of metadata

No one disputes that you need a thick layer of metadata to get computers to perform useful work for you²¹. The more metadata you have, the more a computer can do. What can be in dispute, particularly if you are not aware of the differences between the approaches, is the different ways you can get that thick layer of metadata. Here are the two ways:

²⁰ Introduction to Intelligent Agents for Business Professionals (DRAFT), http://xbrl.site.azurewebsites.net/2016/Library/02_IntroducingIntelligentAgents.pdf

²¹ Understanding the Need for a Framework and Theory, <http://xbrl.squarespace.com/journal/2015/9/20/understanding-the-need-for-a-framework-and-theory.html>

- **Have the computer figure out what the metadata is:** This approach uses artificial intelligence, machine learning, deep learning, neural networks, and other high-tech approaches²² to detecting patterns and figuring out the metadata.
- **Tell the computer what the metadata is:** This approach leverages business domain experts and knowledge engineers to piece together the metadata so that the metadata becomes available.

Both approaches can be appropriate in the right situation. If the tolerance for error is low, like in a financial report, machine learning and neural networks generally don't work. What works really well is when you tell the computer what the metadata is first which serves as training data and then machines can be employed to leverage that training data to derive more metadata which human domain experts sort through and make right.

So again, both approaches can be useful in the right situation. Understanding the pros and cons of each approach helps you make good choices of which approach to employ. Technology has its limits which computer science professionals sometimes forget to mention to business professionals.

1.4.2. Benefits offered by expert systems

In the future, the accounting and reporting rules will exist in both human-readable and machine-readable form and will drive the expert systems and intelligent software agents which professional accountants use to create financial reports. Benefits from the use of expert systems and intelligent software agents include:

- **Automation:** elimination of routine, boring, repetitive, mundane, mechanical tasks that can be automated
- **Consistency:** computers are good at performing repetitive, mechanical tasks without variation whereas humans are not; computers do not make mistakes and are good at repeating exactly the same thing each time
- **Diligence and tenacity:** computers excel at paying attention to detail; they never get bored or overwhelmed and they are always available and will keep doing their job until the task is complete with the same attention to detail
- **Reduced down-time:** computer based expert systems are tireless and do not get distracted
- **Availability:** such computer based expert systems are always available simultaneously in multiple places at one time; you get quick response times and can replace absent or scarce experts
- **Training:** the best practices of the best practitioners can be available to those that are new to and learning about a domain of knowledge
- **Longevity and persistence:** computer based expert systems do not change jobs or retire so knowledge gathered by an organization can remain within that organization
- **Productivity:** computer based expert systems are cheaper than hiring experts and costs can be reduced at the same time that quality increases resulting in increased productivity

²² *Understanding 'Deep Learning' and 'Neural Networks'*,
<http://xbrl.squarespace.com/journal/2015/9/17/understanding-deep-learning-and-neural-networks.html>

- **Multiple opinions:** Systems can integrate the view of multiple experts within a single system and choose between the preferred view of multiple expert opinions in the same system
- **Objectivity:** computers apply the same inductive and deductive logic consistently; emotion and personal preferences can be eliminated where they should be eliminated

Critical to understanding the sorts of tasks that expert systems will be capable of performing and should not or will never be able to perform takes the understanding of a domain professional. While computer based expert systems can effectively automate some work, this does not imply that these systems will automate all work or replace humans. They simply won't because they cannot. Computers are dumb beasts. There is a difference between subjectivity and objectivity; there is a difference between a mechanical task and a task requiring professional judgement. Professional accountants need to understand the difference²³.

1.4.3. Machines helping humans

Machines are good at some things. Humans are good at other things. Understanding the difference is important for a lot of reasons. Machines can be made to perform certain tasks effectively and efficiently.

McKinsey published an article which predicted²⁴:

“Many of the jobs that had once seemed the sole province of humans—including those of pathologists, petroleum geologists, and law clerks—are now being performed by computers.”

Understanding the true capabilities of computer software and humans will help you understand what skills you will want to have in order to make sure you can add value to your organization. Those same skills will help you use intelligent software agents appropriately.

Humans augmented by machine capabilities, much like an electronic calculator enabling a human to do math quicker, will empower knowledge workers who know how to leverage the use of those machines.

1.5. Details of Categories of Intelligent Software Agent Sophistication

In the section *Categories of intelligent software agent sophistication* we summarized the different categories of intelligent software agents and the level of sophistication offered. In this section we want to provide more details on the these different categories.


²³ *Introduction to Knowledge Engineering for Professional Accountants*, http://xbrlsite.azurewebsites.net/2017/IntelligentDigitalFinancialReporting/Part01_Chapter02.3_KnowledgeEngineeringBasicsForProfessionalAccountants.pdf

²⁴ KcKinsey Quarterly 2014, *Artificial intelligence meets the C-suite*, <http://www.mckinsey.com/business-functions/strategy-and-corporate-finance/our-insights/artificial-intelligence-meets-the-c-suite>

1.5.1. Generic agents

A generic agent²⁵ is anything that perceives an environment through sensors and acts or reacts upon the environment through effectors.

Generic Agent



The following is the basic control loop of a generic agent:

While true

1. observe the world (**environment**);
2. update internal world model (**state**);
3. deliberate about what intention to achieve (**goal**);
4. use means/ends reasoning to get a plan for the intention (**condition-action rules**)
5. execute the plan (**action**)

End while

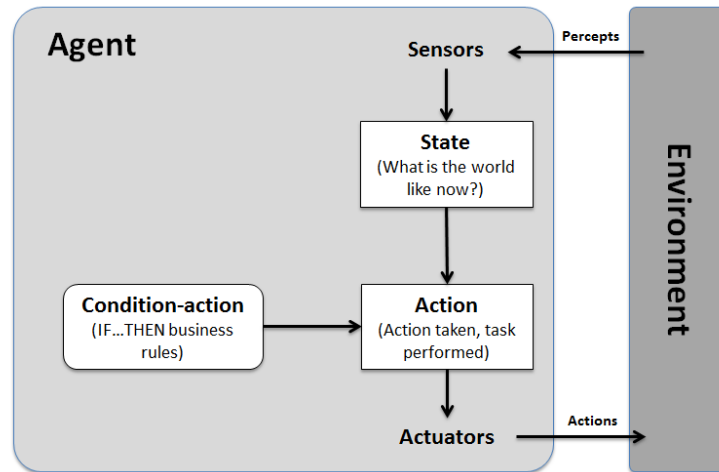
1.5.2. Simple reflex agents

A simple reflex agent²⁶ looks up what it should do from a list of rules in response to its perception to the environment.

²⁵ Consortium of Cognitive Science Instruction, *Introduction to Intelligent Agents, Generic agent*, Retrieved July 24, 2016; http://www.mind.ilstu.edu/curriculum/ants_nasa/intelligent_agents.php

²⁶ Wikipedia, *Intelligent Agents, Simple Reflex Agents*, retrieved August 14, 2016, https://en.wikipedia.org/wiki/Intelligent_agent#Simple_reflex_agents

Simple Reflex Agent



A simple reflex agent looks up what it should do from a **list of rules** in response to its perception to the environment. This is the algorithm of simple reflex agent²⁷:

Function Simple-Reflex-Agent (percept) returns action

- **persistent:** *rules*, a set of condition-action rules
- *state* <<<< Interpret-Input (*percept*)
- *rule* <<<< (Rule-Match (*state*, *rules*))
- *action* <<<< *rule*.Action
- **return** *action*

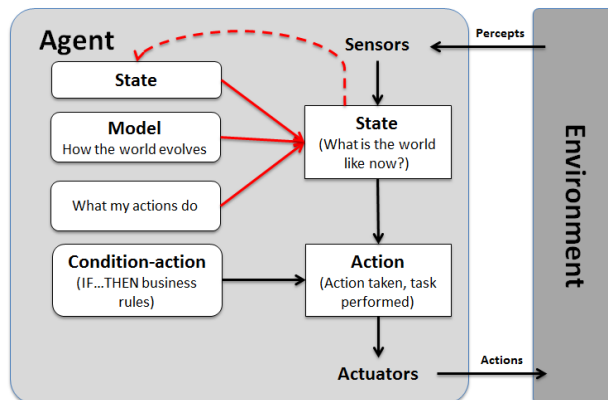
1.5.3. Model-based reflex agents

A model-based reflex agent²⁸ is the same thing as a simple reflex agent except that a model-based reflex agent has a model of how the environment evolves.

²⁷ Hugo Larochelle, *Intelligence Artificielle*, Retrieved July 24, 2016; <https://www.youtube.com/watch?v=TUHAVbaBLIq>

²⁸ Wikipedia, *Intelligent Agents, Model-based Reflex Agents*, retrieved August 14, 2016, https://en.wikipedia.org/wiki/Intelligent_agent#Model-based_reflex_agents

Model-based Reflex Agent

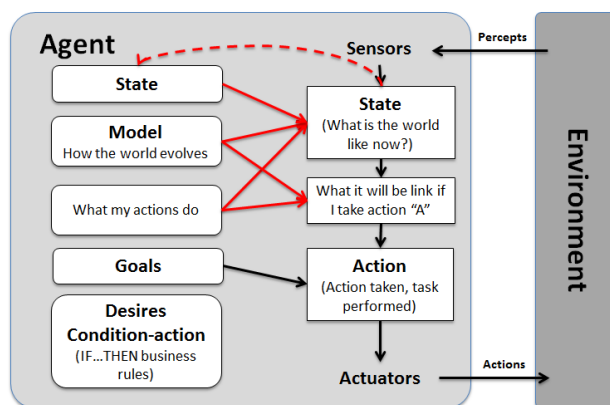


A model-based reflex agent is essentially the same as a simple reflex agent except that a model-based reflex agent adds the additional functionality of a model of the environment in which the agent operates.

1.5.4. Goal-based agents

A goal-based agent²⁹ has a goal or set of goals that it actively pursues in accordance with an agenda (so this type of agent is proactive, not just reactive). A goal based agent has a representation of the current state of the environment and how that environment works. The agent pursues policies or goals that may not be immediately attainable. And so, goal based agents do not live merely in the moment. These agents consider different scenarios before acting on their environments, to see which action will probably attain a goal. This consideration of different scenarios is called search and planning.

Goal-based Agent



A goal-based reflex agent is essentially the same as a model-based reflex agent except that the goal-based reflex agent adds the additional functionality of wondering what the environment will be like if a specific action is taken, evaluating if

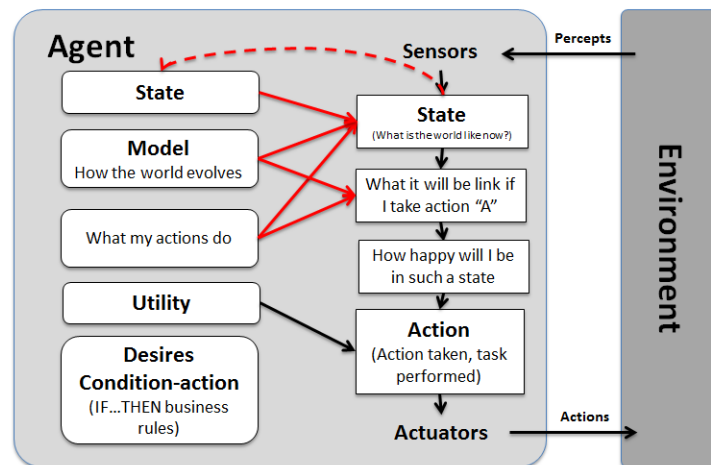
²⁹ Wikipedia, *Intelligent Agents, Goal-based Agents*, retrieved August 14, 2016, https://en.wikipedia.org/wiki/Intelligent_agent#Goal-based_agents

that state is desirable or undesirable given specific goals, and the agent's desires given that state.

1.5.5. Utility-based agents

A utility-based agent³⁰ is a more sophisticated type of goal-based agent that also rates each possible scenario to see how well it achieves certain criteria with regard to production of the good outcome, therefore it is more adaptive. A utility measure is applied to the different possible actions that can be performed in the environment. The utility-based agent will rate each scenario to see how well it achieves certain criteria with regard to the production of a good outcome. Things like the probability of success, the resources needed to execute the scenario, the importance of the goal to be achieved, the time it will take, might all be factored in to the utility function calculations.

Utility-based Agent



A utility-based agent is the same as a goal-based reflex agent but adds additional functionality of wondering what the environment will be like if a specific action is taken, evaluating if that state is desirable or undesirable given specific goals and desires given that state, and evaluating how happy the agent will be within that state. Such agents might even be allowed to change its goals.

³⁰ Wikipedia, *Intelligent Agents, Utility-based Agents*, retrieved August 14, 2016, https://en.wikipedia.org/wiki/Intelligent_agent#Utility-based_agents

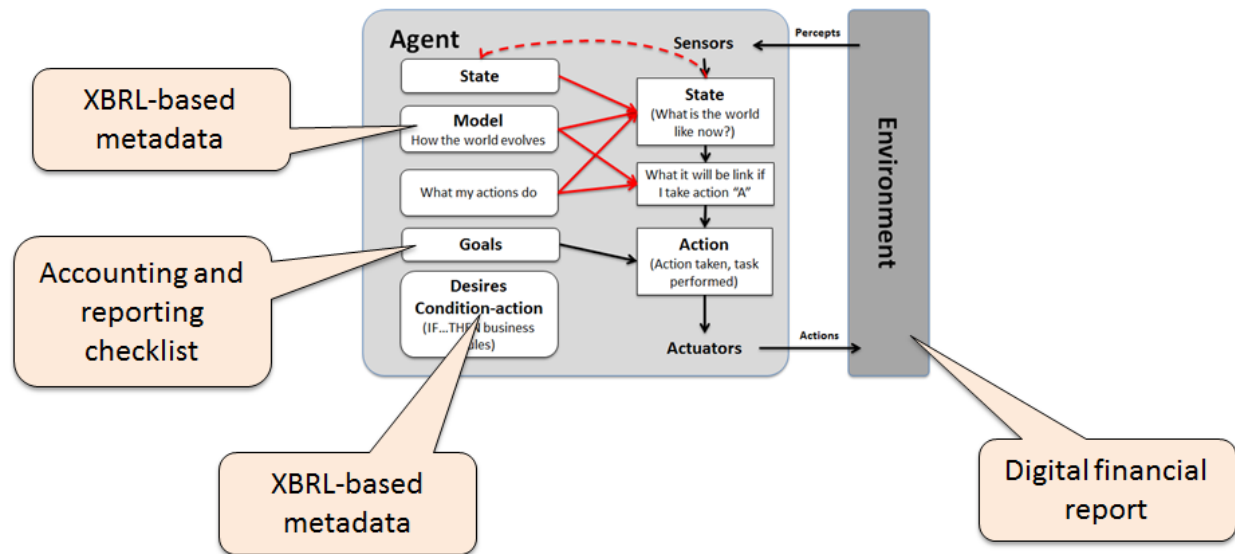
1.6. Using Intelligent Software Agents in Financial Reporting

So what is the best type of agent? That depends. What type of agents could be useful in the process of creating a financial report?

Why would you even want to create an agent for financial reporting³¹?

1.6.1. Starting simple; example of one type of agent

The following is the architecture of an agent that could be helpful in the process of creating a financial report³². I believe that such an agent would be classified as a global standard rational, deliberative, non-learning, goal-based agent.



1.6.2. Process robotics is disrupting accounting and finance

In the video, *Finance in a Digital World and the Impact on CFOs*³³, John Steel who leads the finance transformation practice of Deloitte makes the statement,

"Five years from now there is either no CFO or the CFO is playing a different role."

John's view is consistent with what the AICPA and Journal of Accountancy are saying which is that technology is poised to change the accounting profession³⁴.

In the video, John goes on to say, "Digital is having a tremendous impact and it's quite disruptive." In the video he goes over trends that are occurring. One of those trends is process robotics. John uses the term "lights out finance" meaning a finance

³¹ *Benefits Offered by an Expert System*, Retrieved July 24, 2016; <http://xbrl.squarespace.com/journal/2016/5/30/understanding-the-benefits-offered-by-expert-systems.html>

³² *Automating Accounting and Reporting Checklists*, Retrieved July 24, 2016; <http://xbrl.squarespace.com/journal/2016/5/5/automating-accounting-and-reporting-checklists.html>

³³ SAP, *Finance in a Digital World and the Impact on CFOs*, <http://events.sap.com/sapandasug/en/session/32172>

³⁴ *AICPA News Update: Technologies are poised to reshape the accounting landscape*, <http://xbrl.squarespace.com/journal/2017/7/7/aicpa-news-update-technologies-are-poised-to-reshape-the-acc.html>

that is completely automated. Now, we may never get to where the lights are completely out, but digital will involve automation of many existing manual processes.

A key word here is "disruptive". As pointed out by *The Innovators Dilemma*³⁵, there are two types of innovation: sustaining and disruptive. *Sustaining innovation* meets customer's current needs, making incremental improvements in quality and efficiency of current processes. *Disruptive innovation* is about meeting future needs of customers.

Process robotics is about automating accounting, clerical, administrative, and other such tasks using software robots. Artificial intelligence technology drives these software robots. It really is a lot like how physical robots were employed to replace humans in manufacturing processes such as the process of building cars. Software robots cannot automate all tasks but certainly there are tasks which can be automated.

We don't want to overstate what innovations such as process robotics will provide. But we likewise don't want to understate the impact either. This is not about computers taking over the world. The way it will work is that you will have humans augmented by machine capabilities, much like an electronic calculator enabling a human to do math quicker, will empower professional accountants and others who know how to leverage those machines.

Underlying the process robotics will be intelligent software agents that interact with each other and with humans to get work done.

³⁵ YouTube, *The Innovator's Dilemma*, <https://www.youtube.com/watch?v=yUAtIQDllo8>