



GEORGETOWN UNIVERSITY

Empowering the Knowledge Worker with Artificial Intelligence (AI)

Deploying IBM Watson at MITRE

By

Maryam Sajjad and Travis Johnson

| | |
|---|-----------|
| Abstract | 2 |
| Problem Statement | 3 |
| Project Scope | 3 |
| Research and Data Analysis | 5 |
| Technical Approach | 8 |
| Key Risk Factors | 32 |
| Communication Plan | 11 |
| Change Management Procedures | 12 |
| Solution Development | 12 |
| Technical Architecture | 17 |
| Enterprise Data Governance | 19 |
| Cyber Risk Assessment | 19 |
| Competitive Analysis | 20 |
| Proposed Solution | 22 |
| Business Case and Financial Analysis | 25 |
| Financial Analysis | 27 |
| Ethics | 30 |
| Conclusion | 31 |
| References | 32 |

Abstract

In the past few years, major advancements in the artificial intelligence (AI) sector have peaked the interest of many companies. In 2019, experts see growing confidence that this smart, predictive technology, bolstered by learnings has picked up in its initial deployments and can be rolled out at a wholesale level across all business operations (Marr, 2019). Many businesses are asking if this is the right time to invest in the technology, and one of them is the MITRE corporation.

Our proposal will explore the feasibility of introducing IBM Watson, an AI tool, into the MITRE organization. We will conduct a full analysis of the following areas: Technical, Cyber Risk Analysis, Ethical, Finance and Business Case. This paper will highlight pros and cons of investing in IBM Watson, so that MITRE can make a decision accordingly.

Problem Statement

The MITRE Corporation's mission-driven team is dedicated to solving problems for a safer world. MITRE has approximately 8,600 employees in 70 plus sites around the world. They are a not-for-profit company that operates multiple federally funded research and development centers (FFRDCs). They work across the government, through FFRDCs and public-private partnerships, to tackle problems that challenge our nation's safety, stability and well-being. Their unique vantage point allows them to provide innovative and practical solutions in various spheres, including defense and intelligence, aviation, civil systems, homeland security, judiciary, healthcare, and cybersecurity. In 2018, MITRE's gross revenue topped 2.6 billion dollars. MITRE employees are primarily knowledge leaders and workers. A knowledge worker is defined as those whose roles require both experience and the ability to collect and analyze data needed for decision making and action.

Our proposal will explore the feasibility of introducing an artificial intelligence product into the MITRE organization, which focuses on delivering in-depth research that leads to new ideas for solving problems facing governments at the local, national and global level. We will conduct a thorough analysis on the cost, product maturity, technology and infrastructure needed. Our research will determine if it is the right time for MITRE to invest in this technology.

Project Scope

Artificial intelligence is a theory and the development of a computer system that normally requires human intelligence. A knowledge worker spends countless hours working to understand, research and formulate solutions to highly complex problems. As Figure 1 states, many believe

that AI is ready to assist them by providing the tools and capabilities necessary to dramatically improve research time and accuracy.



Figure 1: Knowledge workers want to eliminate repetitive task (INCP, 2019)

There is ample research and information available that explores the current environment, advances in the field and the pros and cons of this technology. Additionally, there is significant research into the ethical ramifications of AI within the workplace and throughout the world. Our research will discuss these ethical points.

The scope of this project includes implementing an AI product, IBM's Watson, into the infrastructure of the MITRE Corporation. Furthermore, this project will include employee interview findings, implementation needs that include license fees, maintenance cost, hardware and software needs.

Research and Data Analysis

The definition of artificial intelligence (AI) is difficult to describe as the focus of AI shifts depending on the entity that provides the definition. Currently there are at least 6 definitions of artificial intelligence and the research will see how some of the industry's leaders are focusing their efforts (Marr, 2018). Recently, industry leaders began to shift definitions based upon the goals that are trying to be achieved with an AI system. Generally, the industry invests in for one of these three objectives (Hammond, 2015):

The first objective is to build systems that think exactly like humans do ("strong AI"). In 1950, the computing pioneer Alan Turing proposed a test called the The Imitation Game. He proposed that a blind controlled test of verbal indistinguishability could serve that purpose. If a computer program were indistinguishable from people in a kind of open-ended typewritten back-and-forth, the program would have passed the test and, in Turing's view, would merit attribution of thinking (Shieber, 2004). As Figure 2 demonstrates, the Turing Test has been the definitive test to answer the question of artificial intelligence. There have been other tests proposed by scholars, but Turing's have withstood the test of time.

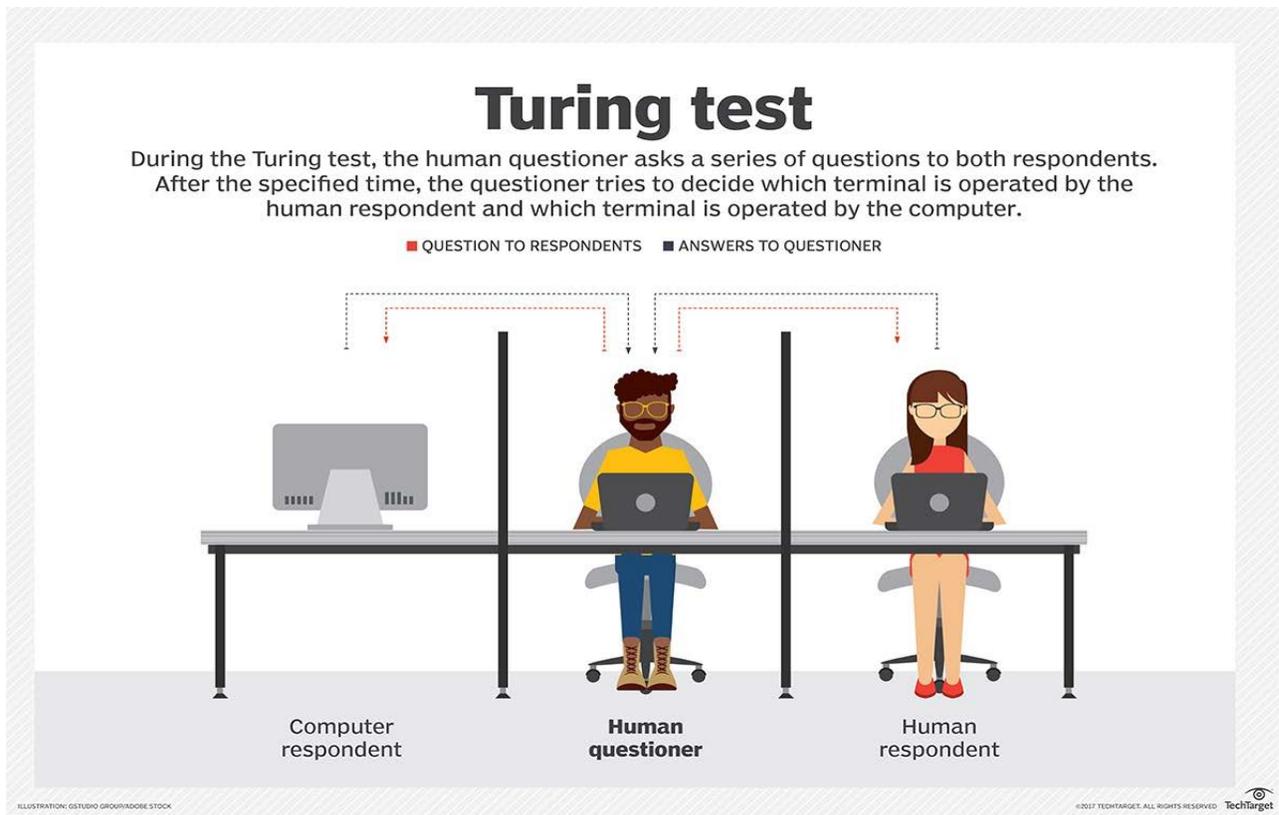


Figure 2: Illustration of the Turing Test (Nguyen, 2018)

This objective not only builds systems that think, but will be able to explain the thought processes behind it, like humans do today. Scientists and researchers have not developed a real model or system that has actual simulations of human cognition.

The second objective is to just get systems to work without figuring out how human reasoning works (“weak AI”). Weak AI seeks to build information-processing machines that appear to have the full mental repertoire of human persons, meaning that they can mimic human reactions and thoughts (Searle , N.D). “Weak AI” can also be defined as focused more on building tools for helping humans in their work rather than on replacing them. As research

suggests, building machines that help humans is a laudable and important enterprise and motivates much excellent AI research (Nilsson, 2005).

The third objective is to use human reasoning as a model, but not necessarily the end goal. This is where most of the AI development is occurring today. Most companies are developing projects that use human reasoning as a guide to creating better products and services. They are not interested in recreating the human mind. The project seeks to arm the knowledge worker through this objective. The goal will be to use AI to assist the knowledge worker with repetitive tasks that include references to a wide range of data resources that reside across the internet, internal storage locations, spreadsheets and databases. AI will be able to compile this data in a meaningful and structured way. Second, AI will be able to assist in providing compelling data visualizations that will help the knowledge worker deliver persuasive information to senior leadership. This way, senior leadership will be able to make better and more informed decisions. Figure 3 demonstrates the augmented AI approach that the objective is designed to achieve.

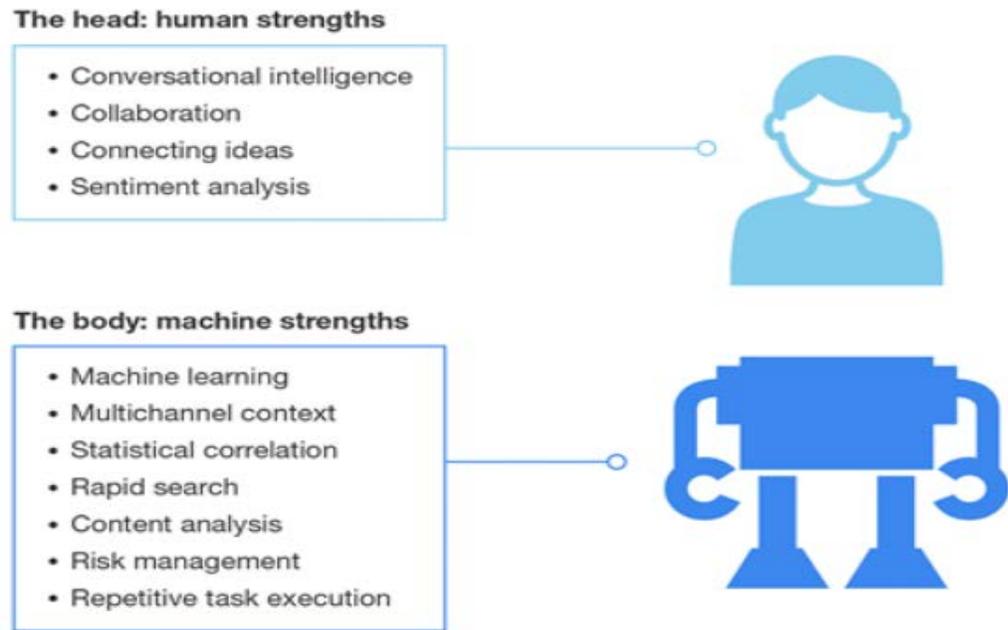


Figure 3: Demonstrating partnership of human and AI (Forrester, 2018)

Technical Approach

Introducing new technologies in an established work environment can be challenging. Projects need to follow an established process that will guarantee a smooth implementation. This project will leverage the preexisting deployment processes of MITRE. The project team will implement this new technology in a phased approach using Agile Scrum Methodology and figure four demonstrates the industry scrum process. The project team will use this phased approach and will be divided by user groups starting with the main MITRE campus located in McLean, VA. The first group will provide the solution team with valuable feedback to make adjustments. The second phase will take place at MITRE's headquarters in Bedford, MA and then the final phases will consist of a controlled roll out to the remaining FRDC sites.

The timeline to implement the new tool, transfer or recreate documents, train the workforce and document lessons learned for other campuses is six to eight months. The

implementation approach will consist of 5 sprints. The first sprint will be used to understand the business requirements and will be labeled at Sprint 0 Discovery. The project team needs to understand the data sources, transformation areas required, and key performance indicators.

Once the project team completes Sprint 0 Discovery, the project team will begin to iteratively work on designing, analyzing, developer testing, quality assurance testing, and measuring the results. Sprints will last no more than 30 days to build a releasable product in every Sprint (James & Walter, 2017). Every day at the same time and place, the project team will spend a total of 15 minutes to inspect the progress towards the sprint goal (James & Walter, 2017). Our focus is to have a limited production release pilot during month 5 of the deployment timeline. While the project team is working to complete each Sprint, training the users and promotion of the solution throughout MITRE will occur.

Scrum Process

Enter your subhead line here



Figure 4: Illustrating the Scrum Process (Brainhub, 2018)

Monitoring and controlling of the process is intended to occur continuously while the work is being executed. To determine the risk factors related to this implementation, the project

manager and other team members will analyze whether the risks are within control or outside of their control, and have a remediation plan ready to respond to the identified risks. Having a risk management plan in place will ensure risks are managed properly, reduce the impact of negative risks, and increase the impact of opportunities (Risk Management Planning, 2016)

For this project, we foresee five major risks that could affect the project. The first risk identified is schedule delays. We want to make sure the project manager is identified in the beginning stages of the project. The project manager and team members will work with the functional managers and other subject matter experts throughout each stage of the project. When the team is heavily involved in planning and estimating all activities involved, schedule delays will be limited.

The second risk identified is requirements creep. As a project progresses, users will request more and more features that were not identified at the beginning of the project. To overcome this, we will involve functional managers and project sponsors to have the needed trade-off discussions, or increase the scope and cost of the project.

The third risk is user adoption. Involving SMEs in various areas will increase the likelihood that users will welcome a new product. Involvement of managers and project sponsors to promote the new tool as intuitive, efficient, and task-enabling will also improve the probability of user acceptance. In addition to involving SMEs, managers, and sponsors, User Training needs to be prioritized to eliminate and reduce user's reluctance to using the new product.

The fourth risk is system integration. During the design and development phases, the project manager and team members need to identify and assess any integration and interoperability issues that could occur. There are physical, logical, and environmental aspects to consider as well as product dependencies with other programs and organizations. Additionally, enterprise

constraints need to be reviewed and addressed. Certain functionality requests may not be feasible due to restrictions expressed in best practices documents and security standards documents.

The fifth and final risk identified is quality assurance. Testing and evaluation needs to occur after each implementation milestone. Testing involves properties such as reliability, security, and interoperability. This strategy provides the project manager and team members with knowledge to measure progress, identify problems, and recommend changes when they are warranted.

Cyber Security Process

MITRE supports many different organizations that require all data to be protected. The solutions team will ensure that the final solution takes into account the various data protection guidelines set forth by our sponsors. The solutions team will have an information security representative throughout the project to make sure all concerns are identified and addressed.

Communication Plan

All five risks listed above have one common theme that can reduce the probability of the identified risks. The common theme across the board is communication. When the project team is communicating with all the identified managers, SME's, sponsors, and other stakeholders, the probability of risk will decrease. This is also why a communication plan needs to be put in place to define the appropriate communication channels. Within the communication plan, stakeholders need to be identified, a communication method for each stakeholder needs to be determined, the frequency of communication needs to be determined, and a responsible owner needs to be assigned.

Change Management Procedures

When the scope, budget, or requirements of the project change, the project development team must have the framework to ensure they can cope with these changes. Even minor changes can affect multiple areas of a project (Turban, Volonino, & Wood, 2015). Initial tasks and requirements need to be identified during the Sprint Planning Meeting. No changes will be made during the Sprint that would endanger the Sprint Goal unless approved by the Change Review Board.

The change management activities will consist of logging the requests for change, categorizing the change requests, assessing the request and the impact of the request, approving or denying the request, scheduling the change, implementing the change, and then testing the change. To avoid delays in meeting the Sprint execution timeline, bi-weekly meetings will be scheduled to keep track of the change request log, the execution timeline, and the budget.

Solution Development

As with any new technology implementation, there are barriers and concerns that must be addressed. One of the major concerns with adopting new technologies is the belief that it will take over jobs or jobs will eventually be obsolete. The solutions development team that will be initiated will include both technical and non-technical personnel. This group will include an executive champion, project manager, subject matter experts, users, a member of information security and a member of the finance department. This group will be engaged throughout the entire product to make sure the project is successful and delivered on time.

| Roles | Responsibilities |
|-------------------------------------|--|
| Executive Champion | Available to ensure engagement throughout the organization and champion the project. Will work closely with peers to maintain vision and organizational support. |
| Financial Analyst | Makes sure the cost assumptions are accurate and that the team stays within budget. If additional funding is necessary, will work with appropriate team members to secure. |
| Project Manager | Manages the entire project from the beginning until closure. They will coordinate schedules, timelines, budget and communications. Works closely with the executive sponsor. |
| Subject Matter Experts (SME) | Assist in determining the costing model, input into the selection process and provide expertise with the implementation and integration. |

| | |
|---------------------------------------|---|
| Users | Assist in the requirements gathering and data collection. Recommends improvements throughout the project. |
| Information Security (InfoSec) | Assist in identifying and addressing security concerns throughout the selection and implementation process. Recommends remediation and assigns security risk. |
| Privacy Teams | Assist in determining potential areas of PII data exposure and works with the legal department to reduce risk. |

The MITRE solutions team will leverage the previously identified option that will use AI to assist the knowledge worker with repetitive tasks that include references to a wide range of data resources that reside across the internet, internal storage locations, spreadsheets and databases. It will be able to compile this data in a meaningful and structured way. Second, AI will be able to assist in providing compelling data visualizations that will help the knowledge worker deliver persuasive information to senior leadership and their sponsors. Senior leadership and the sponsors will be able to make better and more informed decisions. This type of intelligence is where normal human intelligence is supplemented through the use of technology in order to help people become faster and more accurate at the tasks that they're performing (Luxton, 2019)

The key to getting the most out of augmented AI is having a “data-friendly ecosystem with unified standards and cross-platform sharing.” AI depends on data that can be analyzed in real time and brought to bear on concrete problems. Having data that are “accessible for exploration” in the research community is a prerequisite for successful AI development (West & Allen, 2018). The data required for this project will come from research into IBM Watson product offerings, infrastructure requirements for MITRE, data sources ingest into the product (both within and outside the company) and potential security concerns.

The data analysis approach will be threefold. First, is to determine the needs of employees by conducting structured questionnaires, soliciting information that includes, current data sources used, frequency of use, delivery methods and open-ended suggestions. These findings will be used to determine which data is structured, meaning clearly defined data types that is easily searchable. As Figure 5 outlines, examples of structured data include relationship databases like airline reservation systems, inventory control, sales transactions, and ATM activity. Unstructured data which is that does not have a predefined schema or data models. Examples of this type of data are emails, social media, websites and text messages.

| | Structured Data | Unstructured Data |
|-----------------------------|---|---|
| Characteristics | <ul style="list-style-type: none"> • Pre-defined data models • Usually text only • Easy to search | <ul style="list-style-type: none"> • No pre-defined data model • May be text, images, sound, video or other formats • Difficult to search |
| Resides in | <ul style="list-style-type: none"> • Relational databases • Data warehouses | <ul style="list-style-type: none"> • Applications • NoSQL databases • Data warehouses • Data lakes |
| Generated by | Humans or machines | Humans or machines |
| Typical applications | <ul style="list-style-type: none"> • Airline reservation systems • Inventory control • CRM systems • ERP systems | <ul style="list-style-type: none"> • Word processing • Presentation software • Email clients • Tools for viewing or editing media |
| Examples | <ul style="list-style-type: none"> • Dates • Phone numbers • Social security numbers • Credit card numbers • Customer names • Addresses • Product names and numbers • Transaction information | <ul style="list-style-type: none"> • Text files • Reports • Email messages • Audio files • Video files • Images • Surveillance imagery |

Figure 5: Structured and Unstructured Data (Datamation, 2018)

The employee surveys identified multiple sources of data to be used by the AI solution. These sources included email, text messages, scholarly websites that include IEEE, Gardener, college repositories, internal databases and knowledge bases. Second, most AI solutions have selections of application programming interfaces (APIs) which are a set of functions and procedures allowing the creation of applications that access the features or data of an operating system, application, or other service. The identified developers will need to determine if relevant APIs exist in the AI catalog or is there a need to design and develop new ones. Once the APIs have been established, the team will need to train the AI product to search the connected data and deliver relevant results, as outlined in Figure 6. This is accomplished by uploading common

MITRE phases, pictures, text, research terms and keywords into the AI system. Lastly, the data analysis findings will be used to develop the technical, financial and business cases.

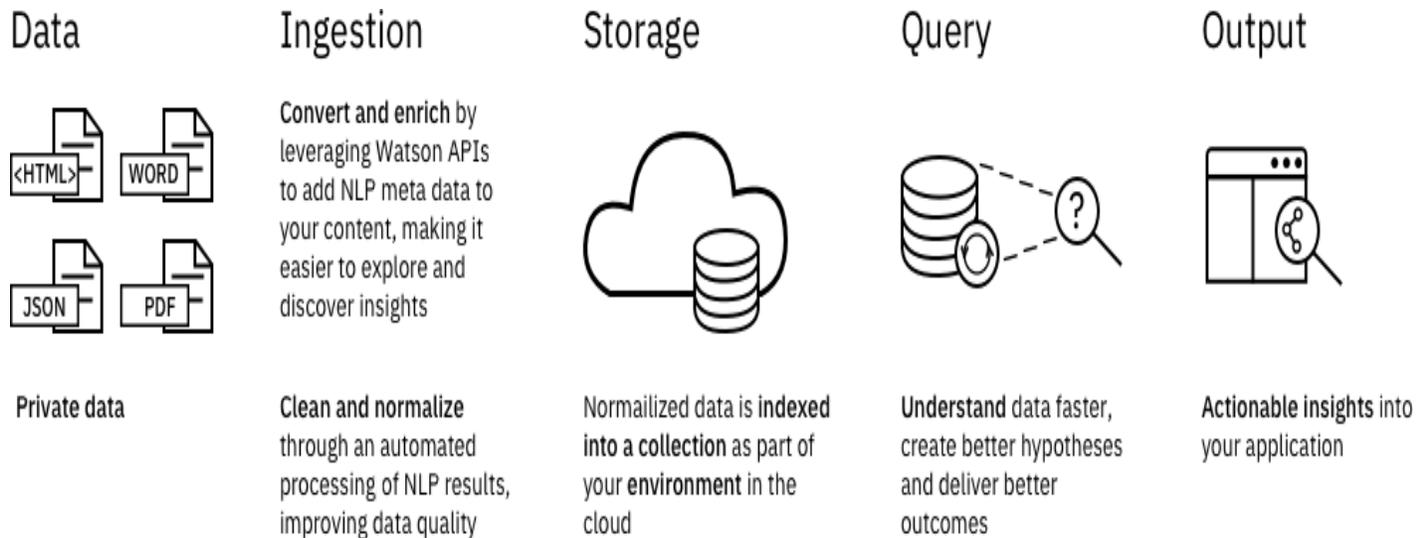


Figure 6: Illustrates Training an AI product (Olavsrud, 2017)

The project will take three years to complete and will be divided into three phases. The first phase will be the discovery phase. During this phase, there will be interviews, market research and a request for information (RFI) sent to select companies. Additionally, a review of current data storage methods and the created of a data governance council will occur. The second phase will use the information collected in phase one to develop a business and financial to present to the leadership team. Once the solutions team has secured the project approval and funding the last phase of the project will commence. The last phase will include the request for proposals (RFP), product selection, implementation and developing key performance indicators (KPIs).

Technical Architecture

The technical architecture approach will consist of a universal semantic layer design. As seen in Figure 7, a universal semantic data layer is a single business representation of all

corporate data. It aims to help end users access all corporate data using common business terms via the business intelligence (BI) and analytics tools of their choice (Olavsrud, 2017). Currently, the data is organized for individual BI tools and designed to be point solutions. The universal semantic data layer design aims to collect all the key existing definitions, business metrics and libraries and place them in a single layer to be managed and changed in one place. Furthermore, the design will allow for the ingestion of external data and collections to be parsed.

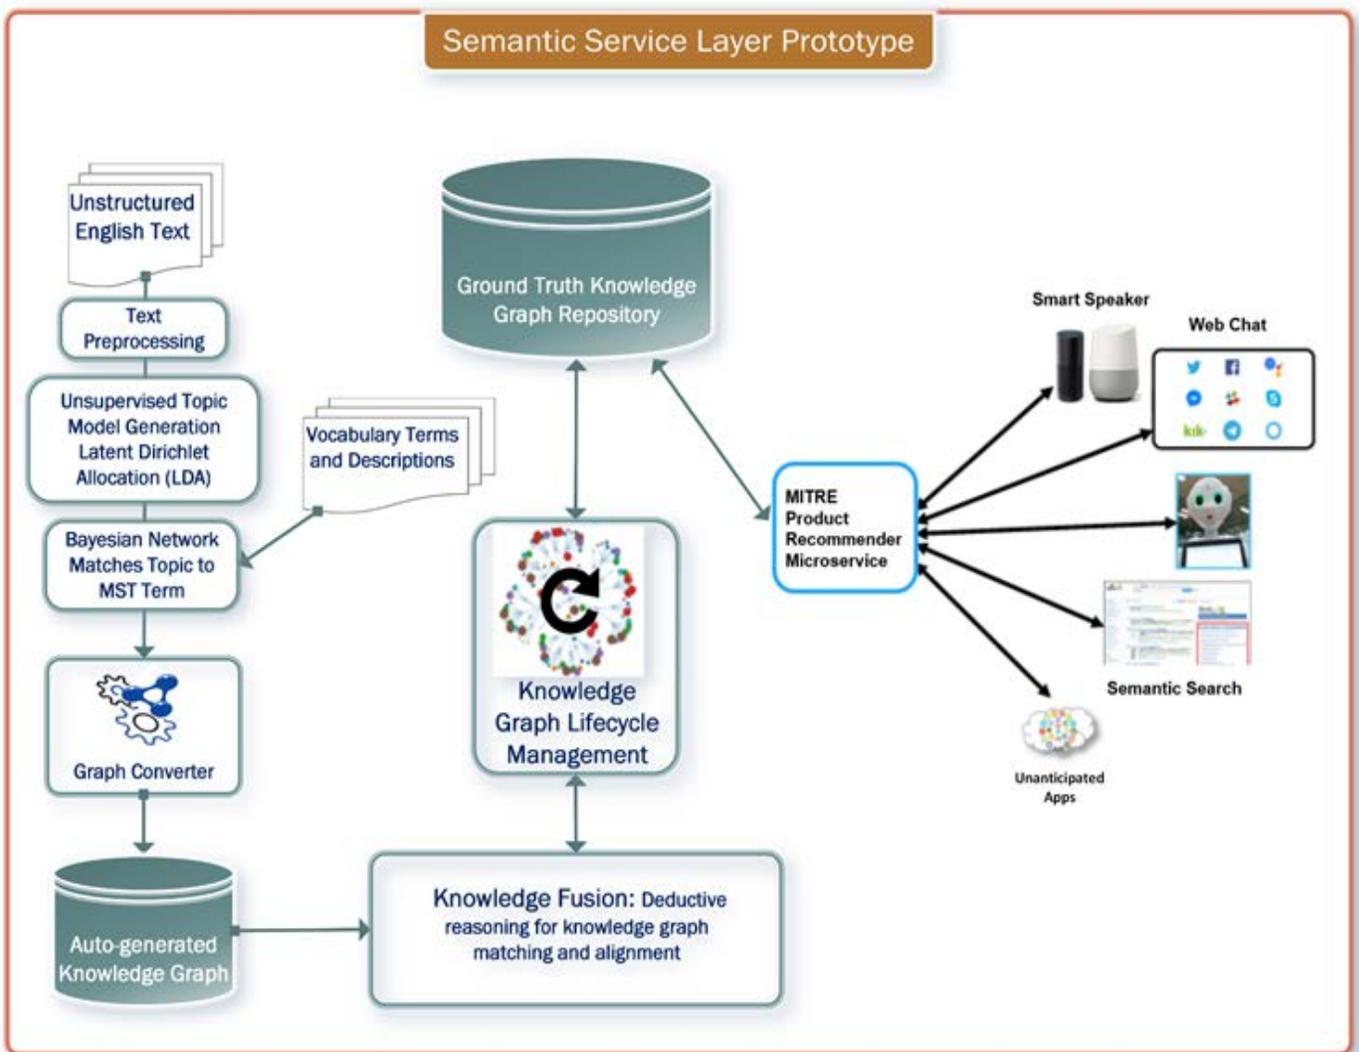


Figure 7: Semantic Service Layer (Olavsrud, 2017)

Enterprise Data Governance

This project intends to establish an enterprise data governance structure that will provide a data governance model that closely aligns with the recommendations of the Data Governance Institute (DGI) framework. The DGI is a logical structure for classifying, organizing, and communicating complex activities involved in making decisions about and acting on enterprise data. The recommended approach is the creation of a new department that has oversight of the newly created data governance office (DGO). This group will have the authority to set and enforce data policies and procedures. In addition, the new members will be selected from across the organization and will have the skills necessary to create, use, set rules and requirements for data.

The new members will collect and report metrics to the stakeholders. They must possess the skills to provide stakeholder communication, access to information and education and support. We recommend that members be comprised of data stewards and stakeholders. This new office will work closely with the developers and product managers to make sure accurate and structured data are maintained in the enterprise.

Cyber Risk Assessment

The solutions team developed the below chart to illustrate the cyber risk they believe could affect the project and the selected product. This chart will be used by the information security team member to discuss any gaps and mitigations vendors may have with their products. The below chart will be evaluated periodically for updates and quality control.

| Risk # | Identified Threat | Impact | Likelihood | Mitigation |
|---------------|--------------------------|---------------|-------------------|---|
| 1 | Unauthorized Access | High | Low | Creation of access controls |
| 2 | Loss of data | High | Low | Regular backups |
| 3 | Data manipulation | High | Low | Access controls and patching |
| 4 | Disruption of service | Medium | Medium | Install redundancy for servers and infrastructure |
| 5 | Process failure | High | Low | Regular training sessions and reviews |
| 6 | Unauthorized data spill | High | Low | Regular training sessions and reviews |

Competitive Analysis

The research indicated that option three, use human reasoning as a model, but not necessarily the end goal, is the selected AI path for MITRE. MITRE needs an AI capability that will assist the knowledge worker by reducing research time and producing data visualizations. The knowledge worker is needed to connect ideas, bring teams together and connect multiple ideas. The research used a combination of market research, SME and employee interviews, cost and security concerns to determine the AI product to deploy. Through the analysis we have

determined to focus the competitive analysis on four products: Amazon AWS, Microsoft Azure, Google Cloud and IBM Watson.

Amazon AWS has the largest offering on the market and were early adopters of an AI platform. Their suite offers a scalable platform that provides a comprehensive data analytics services and powerful predictive capabilities, developer tools such as APIs and software development kits. AWS provides adequate security features at the granular level. AWS does suffer from a confusing billing structure, all data must be stored with them, their application structure is complex, and it takes a lot of time to understand how to use it. The default offering does not include the enterprise version and there have been significant service outages (Opala, 2018).

Microsoft Azure offers the known name branding and has gained a lot of traction in recent years. Their offering includes a scalable model that has a high availability and redundancy model in their global data centers. Their pricing model includes a pay as you go pricing that allows companies to control their budgets. They have a standard security model with strong cybersecurity controls to deliver multi-level protection to end users and business data. Azure needs to be maintained by IT experts that understand patch management and server design and offers less automation options for developers (Opala, 2018).

Google Cloud offers an open source computing environment that can drive down costs, create accessibility, and reduces the requirements for server and infrastructure expertise. Their solution provides multiple of the self APIs and provides Google's state of the art search algorithms. On the other hand, there is no reliable support from google, there is a steep learning curve and you need to store and stage data in their cloud (Opala, 2018).

IBM Watson offers a powerful data ingestion engine and machine learning-as-a-service. The machine learning component is focused on helping developers in getting their models into production. There is access to over 30 types of data stores and a large community support presence. On the other hand, Watson has a limited service offering because they are focused towards building machine learning-based apps through API connections (Opala, 2018).

Proposed Solution

The solutions team has performed a thorough comparative analysis of four technical solutions and used a combination of employee interviews, market research, requirements and security to investigate four top market solutions. Through the interview process it was determined that the users needed a solution that could connect into multiple data sources that used defined business terms, MITRE business jargon, data sets and libraries. The solutions team also considered the overall project objectives in selecting a product that will augment the knowledge worker to provide quicker, succinct research and provide improved data visualizations. Considering these inputs, the solution team investigated IBM, Amazon, Microsoft and Google's offerings.

Comparative Matrix of Solutions

| | | Option 1 | Option 2 | Option 3 | Option 4 |
|-----------------------|--------------------|----------|--------------|-----------|----------|
| | | IBM | Amazon | Microsoft | Google |
| Function | | | | | |
| Integrations | APIs | X | X | X | X |
| | Data storage | X | X | X | X |
| | Data ingestion | X | X | X | X |
| | | | | | |
| Security | FedRamp | X | X | X | X |
| | Defined data paths | X | X | X | X |
| | | | | | |
| Infrastructure | Maintenance Cost | X | UNKNOWN | X | UNKNOWN |
| | License fees | X | UNKNOWN | X | UNKNOWN |
| | Integration cost | X | UNKNOWN | UNKNOWN | UNKNOWN |
| | Staffing | X | X | X | x |
| Performance | Scalability | X | X | X | X |
| | Availability | X | MANY OUTAGES | X | X |

Figure 8: Analysis of Competitors

The solutions team reviewed the four functional areas above to determine the best possible solution for the MITRE Corporation. They considered the documented requirements, cost, security concerns and user interviews when making the decision. It was determined the IBM Watson solution would be the best possible solution. IBM Watson delivers the augmented AI capabilities previously identified along with a robust API offering, easy to understand pricing model, performance and security. Figure 9 illustrates the robust products that IBM Watson has to offer.

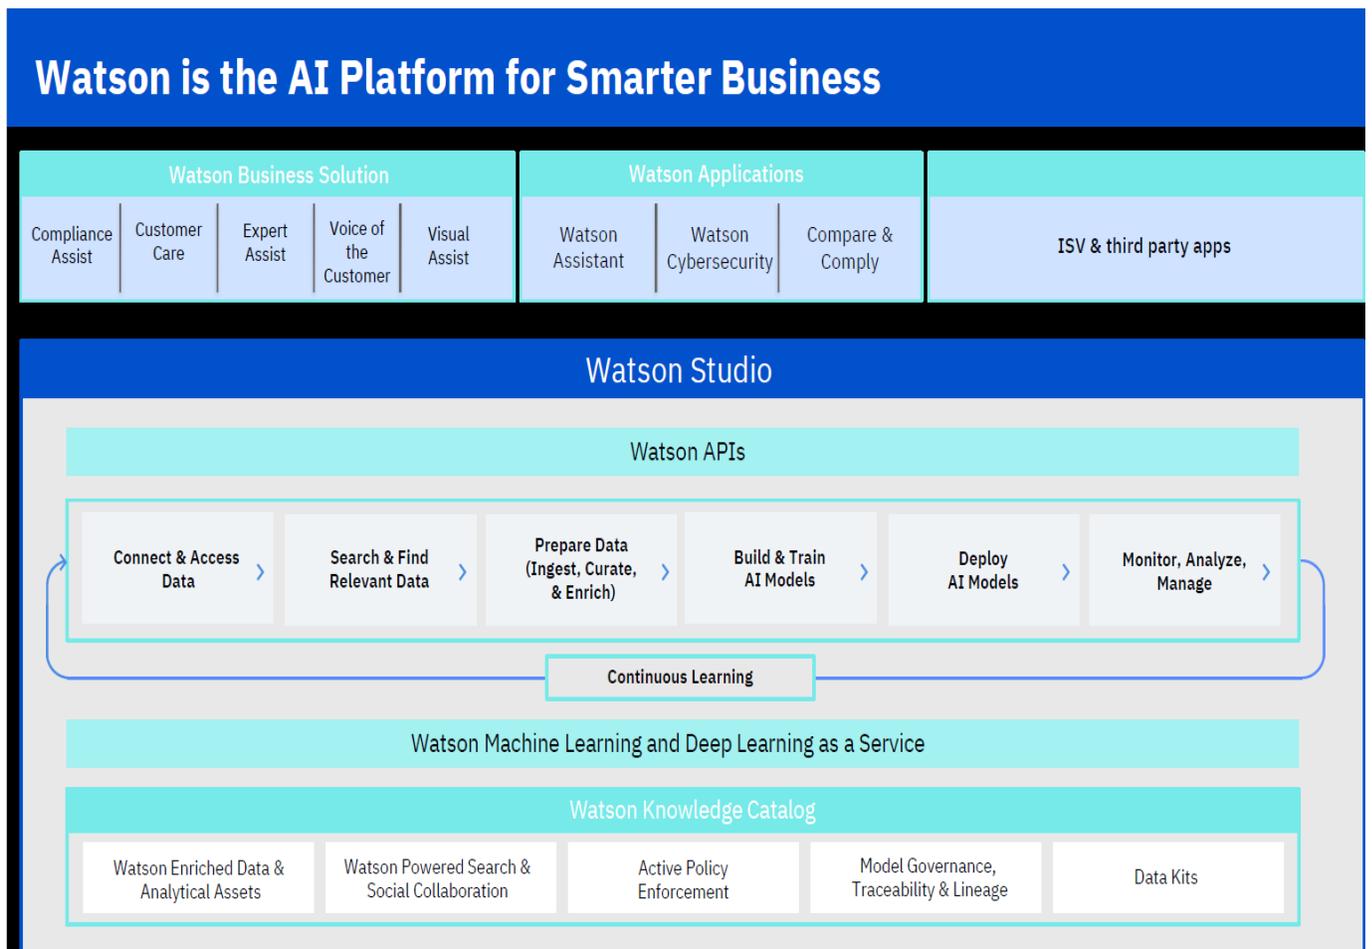


Figure 9: IBM Watson Platform Offerings (IBM, 2019)

The solutions team determined the pricing model for the IBM Watson categories it believes MITRE will utilize in the near future. They will be able to use the numbers in figure 10 to create sound estimated budgets. This Pricing model will allow MITRE to create spending flexibility throughout the year. MITRE will have the ability to increase or limit the services offered by increasing or removing access. For MITRE, this is an invaluable feature that provides unique budget flexibility.

| IBM Watson Services | Freemium | Additional Use |
|--------------------------------------|-----------------------|---|
| Speech to Text | 1K minutes per mo. | \$0.02 per min. |
| - Acoustic Model | N/A | \$0.03 per min. |
| Text to Speech | 1K characters per mo. | \$0.02 per character |
| Watson Assistant | 10K API calls per mo. | \$0.025 per API call |
| Tone Analyzer | 1K API calls per mo. | \$0.0088 per API call |
| Discovery Standard | 2K documents | \$0.0139 USD/hour/1000docs |
| Discovery Advanced (Knowledge Graph) | N/A | \$1K/month \$0.00694 USD/hour/1000docs |

Figure 10: An example of IBM Watson pricing model

Business Case and Financial Analysis

MITRE is uniquely positioned to take advantage of an augmented AI product. MITRE's most important asset is its knowledge workers and providing a technology tool that can provide improved research results, better data visualization models and quicker turnaround for their sponsors, is paramount. There are a few business drivers leading to the decision to investigate and procure an artificial intelligence (AI) product. First, the MITRE sponsors have asked for quicker delivery to problem solutions and advice. In addition, they request better and more sophisticated data visualizations. MITRE's ability to respond quickly and accurately to our sponsor's needs is crucial to the success of the company. Second, the MITRE CEO has recently asked that IT reduce its operating cost, the number of products and applications and improve business analysis.

In the chart below, knowledge workers think that AI is critical in delivering mission critical activities. Over 50% surveyed believe that AI can help with data collection and over 40% think it will improve decision making. In addition, MITRE will use this opportunity to form a data governance office that will create policies and procedures to organize and streamline internal data. This is a crucial step in introducing AI and will reduce the need for multiple business intelligence products.

The MITRE solutions team identified the IBM Watson suite as the option to use to assist the knowledge worker with repetitive tasks that include references to a wide range of data resources that reside across the internet, internal storage locations, spreadsheets and databases. This product will assist them in compiling data in a meaningful and structured way. Second, Watson will assist in providing compelling data visualizations that will deliver persuasive information to senior leadership and sponsors.

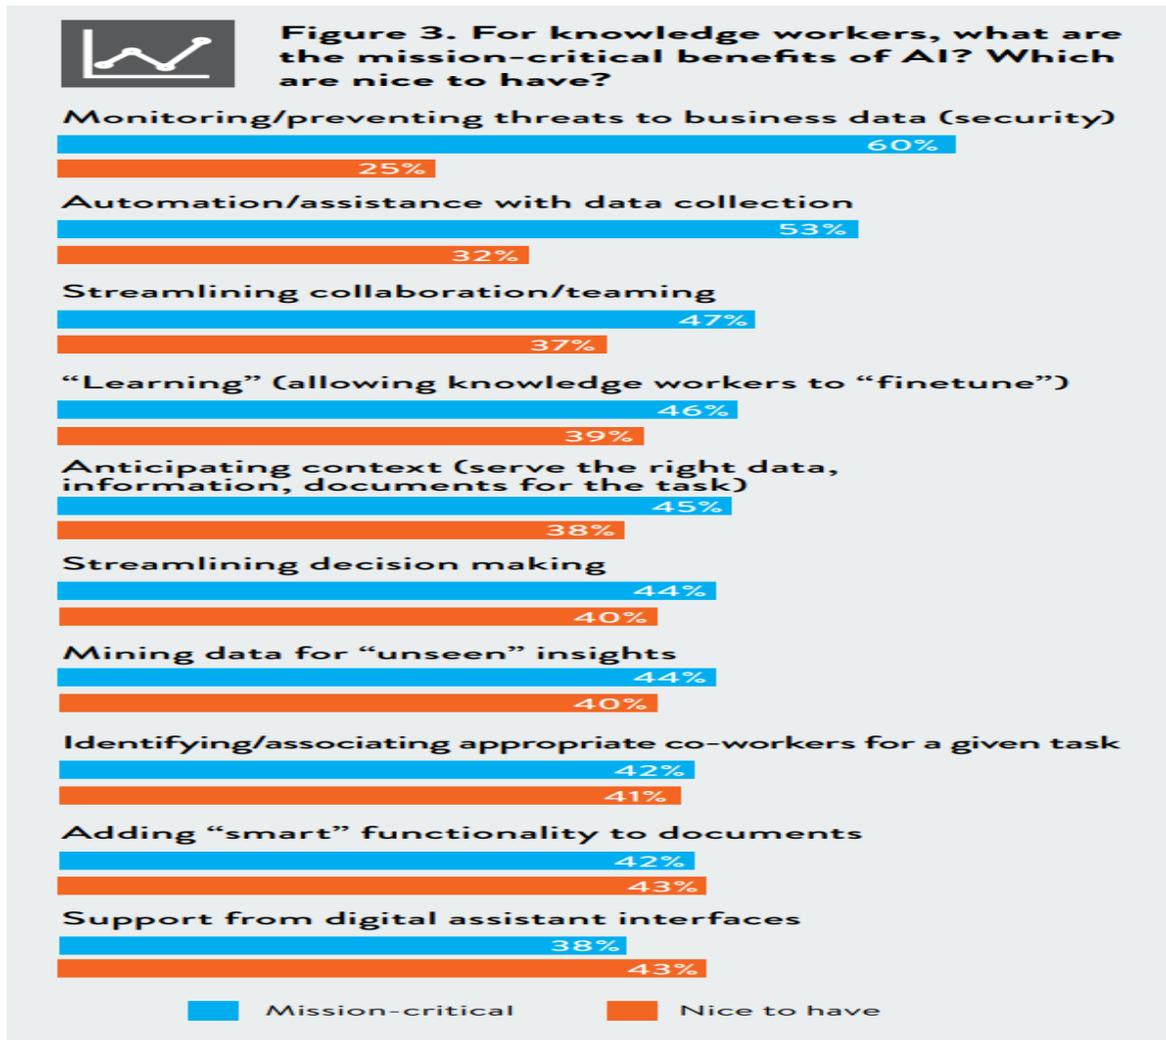


Figure 11: A survey identifying AI uses for knowledge workers (INCP, 2019)

Financial Analysis

MITRE continues to be the thought leader with the Federal Funded Research and Development (FFRDC) groups. This solutions team believes that this project initiative will give MITRE and their knowledge leaders the thought edge in the future. As seen in Figure 12, the financial analysis conducted for this project clearly shows a potential payback of only 2.4 years and a return of invested (ROI) of 72%. There are assumptions that have to be considered for this

analysis. First, the solutions team used the projected IBM pricing chart to determine annual service cost to use the product. The services were selected based on the project requirements and the employee interviews that were conducted. The service cost may fluctuate based on the demand of the product. Additionally, our assumption is that MITRE will have the ability to reduce or add additional IBM services as needed. This will give MITRE the ability to manage the service budget throughout the fiscal year.

The second assumption is the ability for MITRE to reduce or retire established services as the IBM product gains popularity and strong data governance takes root. The assumption is that unused services will be retired in years two and three of the financial plan. This will include all maintenance, license and server fees.

Lastly, the financial analysis assumes that the project will need two developers to aid in designing and implementing the various Application Programming Interfaces which are needed to connect to the data within and outside of the MITRE corporation. The solutions team assumes that these developers will work exclusively on the project for the next five years.

| MITRE - IBM Watson | | | | | | |
|--|------------------|------------------|------------------|------------------|------------------|------------------|
| | Year 0 | Year 1 | Year 2 | Year 3 | Year 4 | Year 5 |
| Revenue | | | | | | |
| Cost Saving | | | | | | |
| Operation Cost Saving (Server, Products, Storage space) | | 150,000 | 250,000 | 300,000 | 300,000 | 150,000 |
| Cost Avoidance | | | | | | |
| Resource cost (2) | | 200,000 | 200,000 | 200,000 | 200,000 | 200,000 |
| Vendor upgrades and maintainiance | | 48,000 | 48,000 | 48,000 | 48,000 | 48,000 |
| Total Cost Saving | | 398,000 | 498,000 | 548,000 | 548,000 | 398,000 |
| Cost | | | | | | |
| Capital Cost | | | | | | |
| Hardware and Software Cost | (150,000) | | | | | |
| Software and License | (130,000) | | | | | |
| Installation Cost | (45,000) | | | | | |
| Project Management Cost | (150,000) | | | | | |
| Professional Services | (40,000) | | | | | |
| CAPEX Depreciation | | (103,000) | (103,000) | (103,000) | (103,000) | (103,000) |
| Total Capital Cost | (515,000) | | | | | |
| Operational Cost | | | | | | |
| Maintenance and 8/5 support | | (30,000) | (30,000) | (30,000) | (30,000) | (30,000) |
| Developers (2) | | (250,000) | (250,000) | (250,000) | (250,000) | (250,000) |
| IBM Watson service | | (70,000) | (70,000) | (70,000) | (70,000) | (70,000) |
| Total Recurring Cost | | (453,000) | (453,000) | (453,000) | (453,000) | (453,000) |
| Total Revenue - Total Cost | | (55,000) | 45,000 | 95,000 | 95,000 | (55,000) |
| Tax Expense | | 19,250 | (15,750) | (33,250) | (33,250) | 19,250 |
| Net Operating Income after Tax | | (35,750) | 29,250 | 61,750 | 61,750 | (35,750) |
| Cash Flow | | 67,250 | 132,250 | 164,750 | 164,750 | 67,250 |
| Cumulative Future Cash Flow | | 67,250 | 199,500 | 364,250 | 529,000 | 596,250 |
| Discount Rate | 2.25 | | | | | |
| Discount Factor | 2 | 2.00 | 1.75 | 1.50 | 1.25 | 1.00 |
| Discounted Cash Flow | (515,000) | 134,500 | 231,438 | 247,125 | 205,938 | 67,250 |
| Present Value of Future Cash Flow | 886,250 | | | | | |
| Net Present Value | 371,250 | | | | | |
| ROI | 72% | | | | | |
| IRR | 22% | | | | | |
| Payback In Years | 2.4 | | | | | |
| Tax Rate | 0.21 | | | | | |

Figure 12: Financial Analysis

Ethics

Ethics should be a core consideration of any action taken by a business. With AI, some of these considerations include building trust, transparency, and preserving data privacy and security.

In regards to building trust, MITRE should consider how AI algorithms are being designed and implemented within the business, as they can potentially discriminate against certain people, beliefs, creed or race. To avoid this, organizations should work towards algorithmic transparency, using control measures to audit algorithms, and doing their best to avoid the “black box” dilemma, which makes it virtually impossible to know the path taken by an AI model to reach a conclusion (BBVA, 2018). We recommend MITRE to place developers under the leadership of ethically sound Managers, that continuously evaluate algorithms to ensure they are aligned with best practices and ethics. Additionally, there should be transparency into the data sources of the information being ingested by AI, so that users can confidently analyze data and be assured of its confidentiality, integrity, and availability.

Apart from the in-house work that MITRE should engage in to protect its ethical standing, we recommend reaching out to professional external organizations like Algorithmic Justice League and Microsoft, which offer guidance on ethical algorithms and dysfunctions of AI. Raising questions to IBM and communicating ethical concerns about their AI product may also be a good idea.

MITRE should also keep in mind that the AI integration is there to enable workers to achieve their goals, not replace them. A special emphasis should be placed on allowing users to work with the technology, so that they can better understand and trust it. As Rob High from IBM mentions, “AI, like most other technology tools, is most effective when it is used to extend

the natural capabilities of humans instead of replacing them. That means that AI and humans are best when they work together and can trust each other” (Middleton, 2017). We recommend the MITRE management to raise a culture that is pro AI at the workplace. A culture that assures employees that although AI will bring positive change when being co-worked with human beings.

As with any technology, AI comes with its fair share of the risk of data being maliciously used or resulting in unintended consequences. As a responsible user of AI, MITRE should consider how to protect themselves under these scenarios, and build their guiding principles and ethical standards accordingly. Having strong cybersecurity practices and measures in place and encouraging overall transparency and trustworthiness will be key in the ethical aspects.

Conclusion

As the research and analysis of this paper suggests, we advise that it is the right time for MITRE to invest in IBM Watson. Through this investment, MITRE will benefit from automating basic processes, enabling knowledge workers to spend less time on formatting, searching and compiling, and allowing them to spend more time on higher, value-adding processes instead (Team Insight, 2019). Additionally, the anticipated ROI of 72% and a payback period of 2.4 years, are major indicators of the financial benefits and desired business drivers. Lastly, MITRE should not overlook the cyber and ethical aspects of this artificial intelligence tool, as that can have detrimental effects. If implemented responsibly, we believe this investment has vast potential for enabling MITRE to achieve its mission and take business to the heights of success.

References

- Bersin, Josh. (2017, March 09) AI, Robotics, and Cognitive Computing Are Changing Business Faster Than You Thought. Retrieved from <https://joshbersin.com/2017/03/ai-robotics-and-cognitive-computing-are-changing-business-faster-than-you-thought>
- BBVA. "How to Make Artificial Intelligence More Ethical and Transparent." *NEWS BBVA*, 16 May 2018, <https://www.bbva.com/en/make-artificial-intelligence-ethical-transparent/>.
- Brainhub. 2018. *Illustrating the Scrum Process*. [Image]. Retrieved from <https://brainhub.eu/blog/differences-lean-agile-scrum/Key Risk Factors>
- Chandra, M. (2016, Dec.). Artificial intelligence and the future of knowledge workers. Retrieved from <https://ieeexplore.ieee.org/document/7784921>
- Datamation. (2018). *Structured vs. Unstructured Data*. [Image]. Retrieved from <https://www.datamation.com/big-data/structured-vs-unstructured-data.html>
- Forrester. 2018. Demonstrating partnership of human and AI. [Slide Image]. Retrieved from <https://www.slideshare.net/nasscom/deep-dive-workshop-i-being-human-in-the-era-of-ai-and-automation>
- Hammond, K., IDG Contributor Network, & Artificial Intelligence Today. (2015, April 10). What is artificial intelligence? Retrieved from <https://www.computerworld.com/article/2906336/what-is-artificial-intelligence.html>
- Henry-Nickie, M. (2018, May 09). AI should worry skilled knowledge workers too. Retrieved from <https://www.brookings.edu/blog/techtank/2017/11/08/ai-should-worry-skilled-knowledge-workers-too/>
- IBM. 2019. IBM Watson Platform Offerings. [Image]. Retrieved from <https://www.ibm.com/watson/developer/>
- INCP. 2019. A Survey Identifying AI uses for Knowledge Workers [Graph]. Retrieved from <https://www.incp.org.co/wp-content/uploads/2019/02/Harnessing-AI-to-Empower-the-Knowledge-Worker-14022019.pdf>
- INCP. 2019. *Knowledge workers want to eliminate repetitive task*. [Image]. Retrieved from <http://www.incp.org.co/wp-content/uploads/2019/02/Harnessing-AI-to-Empower-the-Knowledge-Worker-14022019.pdf>

- James, M. & Walter, J. (2017). Scrum Reference Card. Retrieved from https://www.collab.net/sites/default/files/uploads/CollabNet_scrumreferencecard.pdf
- Le Clair, C., O'Donnell, G., Gownder, J., Sridharan, S., Jacobs, I., McKeon-White, W., Lynch, D. (2018, March 27). AI Is Ready for Employees, Not Just Customers. Retrieved from <https://reprints.forrester.com/#/assets/2/73/RES141543/reports>
- Luxton. (2019). Should Watson Be Consulted for a Second Opinion? *AMA Journal of Ethics*, 21(2), E131-137. doi:10.1001/amajethics.2019.131
- Nguyen. 2018. *Illustration of the Turing Test*. [Image]. Retrieved from <https://medium.com/@vuongptnguyen/evaluating-artificial-intelligence-from-turing-test-to-now-b64a8fced070>
- Nilsson, Nils J. (2005). Human-Level Artificial Intelligence? Be Serious!. Retrieved from <http://ai.stanford.edu/~nilsson/OnlinePubs-Nils/General%20Essays/AIMag26-04-HLAI.pdf>
- Malone, Thomas.(2019, January 24). How AI is changing knowledge work: MIT's Thomas Malone. Retrieved from <https://www.technologyreview.com/s/612815/how-ai-is-changing-knowledge-work-mits-thomas-malone/>
- Marr, B. (2018, February 14). The Key Definitions Of Artificial Intelligence (AI) That Explain Its Importance. Retrieved from <https://www.forbes.com/sites/bernardmarr/2018/02/14/the-key-definitions-of-artificial-intelligence-ai-that-explain-its-importance/#6919336c4f5d>
- Middleton, S., & Stratejos. (2017, January 01). How AI will help knowledge workers. Retrieved from <https://venturebeat.com/2016/12/31/how-ai-will-help-knowledge-workers>
- Olavsrud, T. (2017, November 10). Why your BI strategy needs a universal semantic data layer. Retrieved from <https://www.cio.com/article/3236566/why-your-bi-strategy-needs-a-universal-semantic-data-layer.html>
- Opala, Mateusz. (2018, Feb. 13). Amazon AWS, Microsoft Azure, Google Cloud and IBM

Watson Comparison: How to Choose the Right Cloud-Based Machine Learning Platform for Your Business? Retrieved from <https://www.netguru.com/blog/amazon-aws-microsoft-azure-google-cloud-and-ibm-watson-comparison-how-to-choose-the-right-cloud-based-machine-learning-platform-for-your-business>

Rao, A. (2017, May 10). A Strategist's Guide to Artificial Intelligence. Retrieved from <https://www.strategy-business.com/article/A-Strategists-Guide-to-Artificial-Intelligence?gko=0abb5>

Risk Management Planning – Project Management. (2014, August 15). Retrieved from <https://opentextbc.ca/projectmanagement/chapter/chapter-16-risk-management-planning-project-management/>
Searle, J. R. (1998). *The mystery of consciousness*. London: Granta.

Searle. (n.d.). Artificial Intelligence (Stanford Encyclopedia of Philosophy). Retrieved from <https://plato.stanford.edu/entries/artificial-intelligence/>

Sharkey, N. (2012, June 21). Alan Turing: The experiment that shaped artificial intelligence. Retrieved from <https://www.bbc.com/news/technology-18475646>

Shieber, S. (2018, February 01). Inverting the Turing Test. Retrieved from <https://www.americanscientist.org/article/inverting-the-turing-test>

Team, Insight (2019). Everyday AI: Harnessing Artificial Intelligence to Empower the Knowledge Worker. Retrieved from <https://www.incp.org.co/wp-content/uploads/2019/02/Harnessing-AI-to-Empower-the-Knowledge-Worker-14022019.pdf>

Team, Insight. (2019, February 05). Getting Practical: How AI Is Enabling Knowledge Workers. Retrieved from <https://www.forbes.com/sites/insights-microsoftai/2019/02/04/getting-practical-how-ai-is-enabling-knowledge-workers/#3d88cba81f80>

Tripps, Ready. (2017, Oct. 15). The code of ethics for AI and chatbots that every brand should follow. Retrieved from

<https://www.ibm.com/blogs/watson/2017/10/the-code-of-ethics-for-ai-and-chatbots-that-every-brand-should-follow/>

Turban, E., Volonino, L., Wood, G.R., (2015). Information Technology for Management Digital Strategies for Insight, Action and Sustainable Performance (10th ed.). Hoboken, N.J.

West, D. M. (2018, May 09). Will robots and AI take your job? The economic and political

consequences of automation. Retrieved from <https://www.brookings.edu/blog/techtank/2018/04/18/will-robots-and-ai-take-your-job-the-economic-and-political-consequences-of-automation/>

West, D. M., & Allen, J. R. (2018, May 09). How artificial intelligence is transforming the

world. Retrieved from <https://www.brookings.edu/research/how-artificial-intelligence-is-transforming-the-world/>