



# Augmenting Knowledge Using Artificial Intelligence

How to improve customer (and stakeholder) interactions using artificial intelligence tools to augment the value of knowledge management in human-to-human or human-to-machine interactions.

Author: Esteban Kolsky, **ThinkJar**



THINKJAR  
estebankolsky.com

## VERINT®

# Executive Summary

**We set out to find the best way to use artificial intelligence (AI) to improve knowledge management (KM) in customer service. We comingled a few dozen interviews of practitioners adopting AI for their KM solutions with other research to create two frameworks to build a better system.**

The first one is focused on the five stages of how an AI solution adds value:

1. **Perception.** Determine there is an event that must be solved.
2. **Capture.** Get all the necessary information to frame the problems.
3. **Understand.** A detailed comprehension of the problem or question, and the solution.
4. **Use.** Find the right answer (leveraging available tools), deliver it, and collect feedback.
5. **Learn.** Based on the outcome, identify success and failures to improve future events.

The second one is more focused on how it can add value to an existing KM model, and has four stages:

1. **Index.** Understand the applicable events and create an index to find it when needed.
2. **Storage.** Making it easier to find, by the index, means it needs to be accessible.
3. **Retrieval.** Quickly deliver it to where is needed and ensure distribution and use.
4. **Maintenance.** Based on the usage outcome, improve all items as needed.

Along the way, we used the lessons learned from the interviews, and the many years of research, to provide a basic model of how to generate the right information (using knowledge, content, and data to make sure all aspects of all inquiries are answered), how to make sure that the answer

used is the ONE answer that the customer needs (very seldom done automatically – yet), and set the basic requirements of companies to use AI to augment KM in customer service - while still remaining aware that it is not a panacea that will solve all their ills, but rather a set of tools, a discipline, that can deliver true value.

We conclude the report with a framework that delivers the questions that must be asked to undertake the transition of augmenting KM with AI, and provide a couple of case studies that showcase how BMW financials and Wolseley have adopted these initiatives with great success.

The biggest lesson learned in this research process was that although many organizations are undertaking the initial steps towards letting AI systems improve the results of their human-assisted contact centers (and while those same systems have improved the use of KM in customer service), we are not yet at the point where self-service and automated systems can take over (we discuss why).

At least, not yet. Thus, the concept of augmenting humans, not replacing them, is what this paper focuses on.



# The Real Value of Artificial Intelligence

## Artificial Intelligence (AI) has been deeply bastardized.

Despite decades of academic and practical applications, the discipline of artificial intelligence has lately been used to mean anything from automatic classification of simple concepts to complex black-box machine learning.

In this latest generation - seventh or eighth, depending on whom you ask - we are making AI the first step towards the concept of singularity, or a badly written sci-fi movie where computers become aware of their intelligence to take over the world and enslave humanity. This is far from the real-world applications of AI, and the lab- and academic-centric applications as well.

In between these extremes are some very useful, yet not well used, cases that exemplify where AI can help organizations. The most common example we have found in this new iteration: augmenting humanity. It is not about replacing humans, as science-fiction movies would have us believe, rather about assisting humans to perform at a higher level. The core concept of AI was never to make sentient machines, rather to focus efforts in making machines more useful to humans by doing what they do best - brute-force processing - and letting us to do what we do best - thinking and evolving.

To counter both the singularity movement and the naysayers that believe computers are not intelligent, we set out to investigate the current crop of AI systems and the results across the world of enterprise software. We conducted 18 in-depth practitioners' interviews as well as scoured research and published reports to find working models.

Looking at what the last crop of AI systems left in organizations, we found the most valuable starting point was knowledge management (KM).

We traced the value evolution for AI for KM, from traditional to today's, and the prospects for it. We find the best examples of augmenting knowledge management using AI and deconstruct them to show the progress we have made (and continue to make) in this generation of AI.

Knowledge management and AI have been used interchangeably. Both are very old disciplines in term of technology, dating back five to six decades, and are interdependent - but not the same. Within knowledge management there are a few functions (index, store, retrieve, maintain) that deserve to be understood better as they relate to AI - and implemented with AI more often.

In the latest attempts to do that we have come short due to our inability to find the necessary knowledge at the right time, place, and for the right need. As customers demand answers, organizations must deliver to them THE ONE answer, not a set of potential results which they must detangle to get what they need.

We found that the single, correct, and timely result is knowledge, wisdom, or answers; a list of potential results is educated guesses and passes the burden of finding the single answer to the customer.

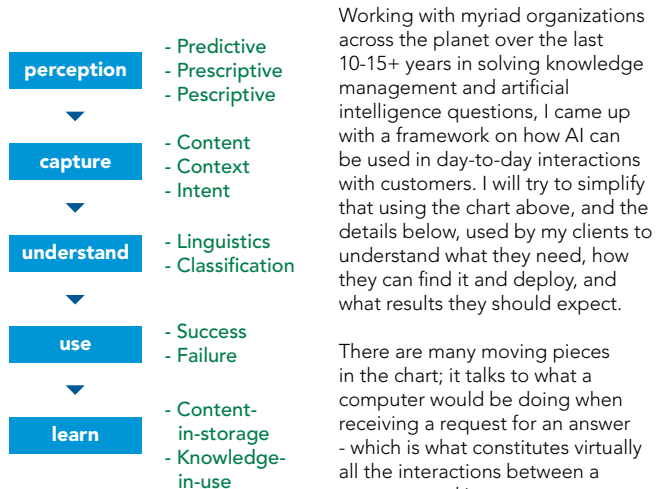
Getting to that one answer has been the domain of knowledge management - albeit executed rather poorly until now.

We will show you how AI today is different from yesterday, and how the right tools and the right focus can make a difference in your organization's adopting AI - and getting good results as opposed to simply doing a trial run, pilot project, or proof of concept of an AI tool - and then staying there.

# The Many Faces of Artificial Intelligence

Without going into excruciating detail, the first Encyclopedia of Artificial Intelligence I read was 3 tomes, each over 1,000 pages. There are many, many books that can tell you what the different tools and technologies used for AI are.

Whether talking about classification, decision-making, rules management, natural language processing, analytics and advanced analytical models, or some other of the commonly-used technologies - that is not what I want to focus on this portion of the paper.



Working with myriad organizations across the planet over the last 10-15+ years in solving knowledge management and artificial intelligence questions, I came up with a framework on how AI can be used in day-to-day interactions with customers. I will try to simplify that using the chart above, and the details below, used by my clients to understand what they need, how they can find it and deploy, and what results they should expect.

There are many moving pieces in the chart; it talks to what a computer would be doing when receiving a request for an answer - which is what constitutes virtually all the interactions between a company and its customers.

## Perception

The start, the genesis of knowledge being used in the organization. There are many ways to perceive (basically, perceiving is to sense some change in state - a customer has a question, there is a problem arising as reported by a machine, an IoT sensor, etc.) but they can be summarized in three modes:

**Descriptive** - there is a description of the need or the question. In more basic terms, a stakeholder (customers, consumer, employees, partners, etc.) asks a specific question.





**Predictive** - a little more complicated, based on learned patterns of perception. Proactive customer service is the most commonly cited example for this.

For example, if an airline flight is canceled, the many customers on it will need to be rebooked. Predictive models use many variables to assist humans in rescheduling those passengers while satisfying the greatest number of variables (frequent flyer, cost of ticket, connections to other flights, return flights are a few of them in this example). The computer suggests a potential solution based on the brute analysis of the many variables for each passenger but does not really know which one will have the best outcome; it simply attempts to recreate past outcomes by correlating variables.

**Prescriptive** - the most complicated of the models; the one that leads to better knowledge and better learning. In this model the computer not only predicts the next-best-action, but also weighs all the different scenarios against each other (and against all that were made available as part of the analysis), and either recommends one over the other based on statistics or chooses automatically (if so entrusted) the best model. In our example, rebooking passengers, the computer defines a platinum loyalty program customer is worth more than a first-time customer – or not, if the data it has is sufficient to make that choice - and recommends

rebooking the loyalty customer over the first-time customer, even if the first-time customer paid more, based on potential future revenue.

In perceiving the data - in some cases before the customer even has access to it - the computer can then assist humans in the organization, or via self-service by themselves, to deal with customers more effectively and with a win-win mentality. Even though the answer won't always be what the customers need or want, it will be the appropriate and correct one, the most logical considering all the available data points and variables in the model.

As complex as it may seem, perception is but the tip of the iceberg of AI. Once we know there is a problem, the real work behind AI commences: capture relevant information, understand it, and create an optimized answer.

Predictive models for KM did not make sense to us until we figured out the secret: it's not about repeating the interaction verbatim, but about having a **"close enough" interaction that could use the same content**

**KM MANAGER**

Global Commercial Bank

Becoming customer-centric changed our approach to collection and use of information; we reduced the amount we captured by almost  $\frac{1}{2}$  and improved first-rate resolution by nearly  $\frac{1}{4}$  by focusing on needed information only.

### VP CUSTOMER SERVICE

Telecommunications

## Capture

Once we, or the computer if allowed to do it automatically, become aware that a potential problem has arisen, we set about collecting the necessary information to solve the problem. There are three parts to capturing the necessary information:

**Content** – the content is the information that encompasses the question or problem: “What is the returns policy for the company?”, “I was overcharged!”, “My product is not working”, or “I want to find the nearest store” are simple examples. Content is decoding the problem or issues into the different elements, so it can be understood and solved.

**Context** – refers to the circumstances surrounding the problem: the reason I need to change my address is because I moved to a new city (does the company serve my new city? Different answers would be applicable under that variable). Context is understanding the variables that make an answer appropriate for a situation.

**Intent** – the most complicated of the three elements, intent refers to understanding the reason the event took place. In some cases, such as a machine that is not working but is necessary for the manufacturing process to continue, the intent is clearer – but when it comes to human interactions, it becomes more convoluted and harder to discern.

If the customer is moving for a job transfer or because they won the lottery, the applicable answer may be different – but without the right information, inferring the intent is near impossible. Intent goes past context, although it relies on it, to focus on the root cause, the origin of the interaction where AI can be applied.

Capturing the right information for the computer to apply its processing power is further complicated by the simple concept that organizations, until not too long ago, thought in a company-centric manner; they did not seem to care much about the customers’ perspective and reasons to do what they were doing.

If a customer wanted to change their address, they would go ahead and do it – without worrying about why or where if it met their standards for processing the change. Once the address was changed, the interaction ended and that was it.

More recently, with the expansion of ubiquitous customer demands we began to look at interactions as customer-centric: revolving around the customers’ needs. Although we have had the concept of customer centricity around for a long time, it was not until social media and extended online and mobile presence that we felt we had enough information to ascertain context and infer intent – and thus were able to understand the customers’ needs and intentions.

## Understand

Understanding is the beginning of the loop of resolution.

Once we understand what the question, problem, or issue is we can begin to resolve it. Until we get to this stage, all we have done is sense there is a problem and capture information about it – but we have not fully understood the problem and all its connotations.

There are two classes of tools that are leveraged by this stage:

**Linguistics** - this is more than likely what you think of when you think of AI today: a set of natural language processing (NLP) tools that can take a statement, a sentence and convert it to a problem and variables associated with it that can be given to a computer in search of a solution.

Linguistics is what most, if not all, organizations, have implemented in the past decade or two as KM for customer service, even for the enterprise.

The idea is that if we can capture the question (including context and intent), we can then associate it with an answer is the basis for linguistics: let the user explain as much as possible what they need, then solve the problem. That's not the way we have done it; we focused almost exclusively on content and simple linguistics tools to match it to search parameters

**Classifiers** - after we understand the language being used, and the associated content, context, and intent based on existing or new data captured, we then need to classify the problem in search of a solution.

To reach a single answer, classifiers are the most valuable tool. As an example, out of potentially hundreds if not thousands of possible answers, classifying the inquiry as a simple change of address, for example, reduced the potential set of answers to dozens; if we can classify it further as an out of town move, maybe a handful; as a mandated move (job related), maybe one or two.

By classifying the problem, we can narrow down the solution – and the faster we can allocate it to the right bucket the easier it is to find the unique answer.

The most interesting part of understanding, despite the fascinating nature of linguistics and classifiers, is the use of a universal descriptor language (UDL). This is where the differentiation between tools begins to take place. An UDL is a personalized language that describes the inquiry and the three parts of it: content, context, and intent. It then creates metadata that makes the inquiry more powerful: information like origin and tracking data, additional information captured, and likely other similar events with resolution.

By using the UDL the question goes from a question asked by a customer to a problem assigned to a computer for resolution. The value of the tools lies in the ability to apply natural language and metadata to a problem, and then translate the problem resolution to a single answer. In other words, UDL translation.



## Use

The simplest aspect of AI and the practical application of the complex prior stages: apply all the collected and understood information and do something. It is also the most structured stage of the AI-driven interaction.

There are two potential results when using AI: success or failure. The structured aspect pertains to the potential success.

For an AI-driven transaction to succeed, there must be an outcome that is correlated to the original stakeholder's needs. If the customer is looking for an answer, or the agent is looking for more information, or an employee is looking for a policy of the company, the process is the same: the outcome is the right answer, at the right time, in the right place.

For this to happen, AI can only deliver an outcome in one of three ways: optimized workflows, personalized solutions, or automated answers

This is where AI systems differ from human-to-human interactions. In a traditional knowledge-aided human-to-human interaction, the answer is given by one person to another after adding a value layer of interpretation.

Computers don't interpret things, not even in AI or machine learning, but rather complete tasks. This is at the core of what AI can do for human interactions, and at the core of how they augment knowledge by learning the best paths to the right answer, and either optimizing the process to get there, personalizing the answer, or automating the interaction completely.

## Learn

Artificial intelligence is the entry point to machine learning. After a machine becomes sufficiently well versed at perceiving a problem, capturing the necessary information to solve it, understanding the inherent questions or issue, and creating the unique solution for it - it begins to learn.

Learning, represented by the science of cognition, is simply an accelerated representation of applied intelligence. Without going deep into cognitive sciences, it's the ability to differentiate between success and failure, and learn how to use the underlying components (in this case information) to repeat successful outcomes.

In the purview of this research, learning is what makes AI augmenting knowledge management possible: taking stored, static content, learning how add value to it by leveraging knowledge, wisdom, and data, and finding the right answer are the entry points to the ultimate outcome of knowledge management (and the leverage of AI tools for it): repeatable outcomes.

In the next section we will elaborate further how this works in the discipline of augmented knowledge.

We worked very, very hard in creating a taxonomy (different categories) but we never used it. Once we began to classify questions by the taxonomy, **our resolution rates shot through the roof.** Like magic.

**VP CUSTOMER EXPERIENCE**  
French Retailer



# Augmented Knowledge as a Discipline

Knowledge seldom is the entire answer. There are three components to generating the right information to get to the right answer: knowledge, content, and data.

Content is static, is “information” that seldom if ever changes (and when it does, is updated by the people who know what’s in it). Content is basic things like manuals, processes, procedures, how-to articles and similar. It’s the address used to send checks, and the phone number to call for support.

Knowledge is the application of that content to a specific situation, bringing to bear the context and the intent of the interaction, as well as ephemeral content (content that is only useful for a very short time or in very specific situations, that resides either in communities or with subject matter experts).

Data, the third component of information, is the operational parameters used to frame an interaction, such as customer number, address from which an email came, time and date when an inquiry came, specific framing data points like SLAs (service level agreements) or contracts and entitlements.


It can also contain very personalized content – uniquely applicable for one interaction or one account, or even a specific client.

The answer for every question comes down to a combination of the three (see chart) and as you can imagine changes every time given the constraints highlighted in this paper: context, intent, time-sensitive needs and demands, and combinations of variables and operational parameters that make every interaction personalized and unique.

Traditional KM used in enterprise is lagging: the focus is in creating content, storing it, indexing it, and retrieving it for later use as necessary.

Alas, the missing components that personalize that content, and that truly create knowledge - content in action - are lost in the store-and-find model of knowledge-in-storage.

Knowledge management augmented by artificial intelligence implementations can generate far better results: unique simple answers to unique personalized problems. This is the sought-after goal for deploying knowledge management: finding THE ONE answer. AI can help you get closer - and eventually, there.





## Index

In the previous section we described the capture and classification of the information stage - and in the discipline of augmented knowledge this is where it happens.

Being able to identify the right information as to where it should be stored, how it should be accessed, how to classify it, and how to use it is traditionally done by a team (in some cases even a single person) of knowledge specialists that tag and index each article or information added to a knowledge base. There are three problems in this model:

1. **Manual, tedious process** - leveraging the tools of AI we can more accurately, and automatically, tag and classify the information in relation to its usage and relative value to each solution.
2. **Slow to react to changes** - it's hard to know when to make changes to the content, the tagging, or the indexing. Information that was useful for a specific use case today may not be the same for tomorrow. AI can pinpoint the necessary changes making it faster to change the content.

**3. Improper maintenance** - while few users provide feedback on content that's outdated or not useful, AI tools can quickly and efficiently show information that may need maintenance thus optimizing the maintenance processes.

Imagine implementing an AI tool that can automatically indicate for the knowledge admins how to tag and index information, not based on existing classifications but based on usage and relations that may not have been known between elements. Or telling them where they should invest their time and efforts on creating new content, tagging existing information better, or even creating knowledge for specific situations that are not being addressed correctly by the system. These promises were made by old KM solutions, but seldom delivered effectively. AI can make them very effective.

All those things are done today in very exhausting processes that involve reading through usage logs, whereas an AI-augmented solution for KM can do all that automatically – even take care of the fix, in certain cases, without needing knowledge administrators.

**Content** - Static, Seldom Changes  
**Knowledge** - Contextualized Content  
**Data** - Operational Parameters

**CONTENT + KNOWLEDGE + DATA  
= INFORMATION**

## Storage

Once the information to be used in knowledge bases is identified, indexed and classified – either by computers or humans – it is stored for future retrieval and use. Although the process sounds simple enough - put an article into a knowledge base - there is more to it and this is where problems arise.

All organizations have multiple repositories of content, data, and knowledge. Finding which of these repositories are useful to whom, whether the formation must be used in different languages, in different regions, which are the most effective and efficient, which has the most flexible models, and other questions that relate to the efficiency of the storage of the information (not to mention that data is usually stored in databases with very different and complex security and rights models) is not simple.

An admin can build a matrix that can then be used to understand better where each piece of information (content, knowledge, and data) can be stored, who should have access to it and how. Optimizing that matrix for any KM solution is a very slow, monotonous, trial-and-error process today. An AI-augmented solution can propose optimization patterns for storage that are far superior to the traditional “just put it in the knowledge base, will figure it out later”.

Imagine an organization that can reduce the size of its knowledge base and all other information repositories, access the information faster, and virtually not worry about what information is valuable to store or use quickly and discard. AI can help with that.

## Retrieval

When a stakeholder asks a question, brings forth an issue, or needs a resolution, they will choose a channel and send a message to the

organization, which will then need to find the right answer utilizing all accessible information and compose the appropriate response.

AI can solve the first part of the equation, understand, given the right parameters and with access to the timely variables. Knowledge management, by contrast, does best the second part of the interaction – finding the appropriate response. The value AI brings is a better formed question for KM to frame a better answer (THE ONE).

If a customer, for example, were to ask “how to balance a portfolio”, a traditional search engine would return both results for how to find the balance and how to allocate the funds among different investments – mostly because it’s unaware of the surrounding circumstances.

We had a knowledge base of tens of thousands of articles, but we were only using a handful of them: agents were filling in holes in the knowledge base with what they knew. We stored less articles, but more relevant, and satisfaction in our self-service solution went up more than we expected, and maintenance costs went way down.

**KM ADMIN**

US Healthcare Insurer

Finding the context for the questions became far more relevant to finding the right answer for us than having a huge knowledge base.

**DIRECTOR OF KM**  
Manufacturing Company

However, if an AI-augmented system were to also submit information related to the customer's stage in life (approaching retirement), and other actions prior to asking the questions (navigating retirement pages within the web site, or looking at retirement communities), and were to ask the customer for more information related to the inquiry (have a conversation about why they need to do that via a chatbot) without human intervention, the inquiry is more precise about how a person about to retire needs to allocate their funds better to last them as long as they want.

In this case, an AI-aided solution found not only the question being asked, but also the variables that allow the KM solution to retrieve the right answer, the one answer, faster and to give the customer what they needed - even if they did not specifically ask for it, although they were thinking it.

AI does not just provide a better-formed question for faster, appropriate retrieval. It can also function as an early indicator for potential problems in the retrieval process by noticing and monitoring the follow-on activities after an answer is given.

If the AI system identifies the customer as unsatisfied, and it happens more than once, then it can tell the indexing system that the answer does not work in that situation and tell the retrieval system to ignore that specific information.

Since most effort is spent in finding and retrieving knowledge to create the right information, the AI solution can provide the most value in helping the search by adding more variables and parameters, and the retrieval system by telling it what works and what doesn't because of its learning processes.

Doing this would not only save money by focusing effort on the right information to index, store, retrieve and maintain. It would also help the organization meet its customers' expectations for speed and efficiency in providing the right answer.

## Maintenance

The most contentious part of KM as implemented today is maintenance. Maintenance does not get the necessary budget, focus, or resource allocations - yet, it's the most important action to keep a KM solution running smoothly. It suffers from three major problems:

1. **It's costly** - as a rule of thumb, costs of KM maintenance are about the costs of deploying the knowledge base for the first year, plus a growth of around 8% average per year. Organizations don't allocate nearly enough to start, nor do they budget for the growth in expense. As a result, most knowledge bases are stale and out of sync, making maintenance even more expensive since it adds the "fix what broke because we didn't do it right" cost.
2. **It's always changing** - an old professor of mine used to say that the value of knowledge decays at the rate of fifty percent per day. While we can argue that today the day has become an hour - or even a minute - given the sharp rise in the pace of information growth, the

statement remains the same: knowledge (that applicability of content to specific situations based on multi-variate problems) is always changing as the variables that make the knowledge valid change.

3. **It's very, very hard** - the complexity of keeping up with knowledge is only mildly understood by those who are not knowledge administrators; knowledge validity is tied to the ability of knowledge and content to be the answer to any question. Not a part of a larger guess that could lead to an answer, but an answer. Finding and keeping that single answer constant, with all variables constantly changing, is a very hard problem - but it is also one that was made for the type of brute processing that AI is best suited to do.

The underfunding of maintenance is the number one reason KM systems failed in the past. Not having updated knowledge available, being impossible to find, or not having it when it should be there are the leading causes of user complaints about using KM systems.

Augmenting KM with AI is very useful to the above three areas, but as AI progresses towards machine learning (as AI is implemented and the system begins to learn) the true value of the augmentation comes down to solving the core issues of maintenance.

And in the processing model of AI we introduced above, the two latter stages, use and learn, power that resolution. Using the right knowledge - or using the wrong knowledge, but knowing why it was wrong - is what can solve most of the maintenance issues: overcrowded knowledge bases begin to evaporate as articles that are irrelevant are highlighted and removed; the most relevant content that applies to specific situations is identified - as well as the gaps it has, so an admin can consolidate the information or approve of the different pieces of content used together; and most importantly, the cost of rogue, brute processing by humans is replaced by cheap, but very effective, AI doing a better job.

Finding the context for the questions became far more relevant to finding the right answer for us than having a huge knowledge base.

**DIRECTOR OF KM**  
Manufacturing Company

It would be very hard for AI to find a better discipline to augment than KM - not only because of the reasons cited above, but because AI is about data, and content, and knowledge, and these are the focus of KM disciplines. We are beginning to see the first versions of new AI-augmented KM appear - starting with Knowledge Automation.





## The Emergence of Knowledge Automation

Starting in the early 2010s the demand for new models for KM was almost deafening. Three trends made the need for new KM felt in organizations: 1) the revolution in customer interactions led by online communities and social media; 2) the sheer amount of knowledge known to exist in communities - and the subsequent rise in Google-initiated searches for answers; 3) and the discovery that existing KM was not only very expensive to maintain, but also not very effective.

The search for new models began to take a toll on (then) existing vendors who went in different directions. Some of them became search and traditional KM vendors and perpetuated the model of knowledge-in-storage, but a few brave souls began to play with AI technologies and concepts and to look for answers elsewhere.

From these few vendors and practitioners, we saw the emergence of knowledge automation. The concept of knowledge automation is an evolution from the old models of store-and-retrieve knowledge. While there is value to that model, in a fast-moving, more context-focused world the model was not delivering. Latency was increasing, differentiation between entries was not achieved, and taxonomies were rendered useless when incorporating the myriad variables that organizations were collecting as part of their Big Data and Customer Engagement initiatives.

A query that was created to use a few, in some cases a simple keyword, inputs is no suited for a complex situation where the organization is trying to find the one answer to the problem. Making acute use of the key characteristics of computers and AI, repetitive processing of information to augment the value of it, a knowledge automation project focuses on pre-fetching the answer as more and more information is available. A good parable of what it does is constant, iterative searches

We came across knowledge automation when we exhausted the ability of our then state of the art KM system to find information quickly. We got **better speed**, but also **amazing accuracy in finding the right answer**.

**KNOWLEDGE AUTOMATION USER**  
(company name withheld at their request)



on the knowledge base with more and more information as it is made available until the right answer is found.

In the case we used above, a person approaching retirement age looking for information on how to better balance a portfolio, would start with the content of the question: how to balance a portfolio and then either answer questions from an automated attendant to focus away from just a balance to an action of balancing, to further refine it to be about retirement balance, and to add data from their account to find one answer that would not only answer the content of their question - but also their intent.

These continuous refinements of the search context are what knowledge automation is all about: automated pre-fetching of information that is continuously optimized and personalized as more data is made available to zero in on the answer that matters.

The value of this method, still in its infancy, is the ability of the AI-augmented solution to process millions and millions of potential variables to focus on the similar case to the one being resolved, the ability of the KM systems to find further refined knowledge in record times, and the ability of the knowledge automation solution to put both together. And while issues remain on adoption at scale, the few practitioners that I spoke with about this were absolutely convinced that this is the future of KM and AI working together.





# Seven Lessons Learned

Throughout this project, and the many interviews and conversations about the topic, we came across some phenomenal experiences. We sought not just to tell their story, although we do have a couple of stories below, but more importantly read through all the notes and find common points to highlight.

The lessons we learned while interviewing practitioners and discussing the concepts were simple but very effective. As it relates to using KM aided by AI, use the following seven lessons to frame your initiatives:

1. **AI is about data, KM is about knowledge.** Neither one works independently, and neither can replace the other. You must continue your KM efforts, augmented by AI, to manage content and you must learn to use AI to manage data and the variables that affect the usage of that content in context – which is what knowledge is.
2. **Takes time to do it right.** Don't be fooled by the hype and the many vendors' promises of "quick results". While you may be able to solve a simple problem by adding AI to it, it will be short lived. Just like with KM, proper initiatives with long-term planning and maintenance (or training in the case of AI) are necessary before results become sticky.
3. **There is no one-solution-solves-all-your-problems approach.** This is one of the largest gaps between technology promises and reality – while one out-of-the-box solution may be easy to deploy under certain controlled conditions, the use of AI to augment KM requires extensive planning and training and testing.
4. **Know your outcomes.** The success of AI is measured by ensuring that results are better than they were in expected outcomes. There are three potential outcomes for AI: optimization, personalization, and automation; they are different and vary in complexity and timing to get right. Know which outcome you are going after, and plan accordingly.

5. **Don't be fooled by the past.** You may have done many AI-like or AI-related things in the past: natural language processing, fuzzy logic, decision systems, or others. The new AI, within the framework introduced in section one above, requires using the lessons learned while doing that, and extending it to a full-enterprise, end-to-end leveraged model to deploy AI to specifically augment KM. Don't forget what you learned.

6. **Lather, rinse, repeat.** It bears repeating: the results intended and expected the first time an AI initiative is launched are unlikely to either be obtained or be the main reason for adopting AI. Takes several iterations to align the understanding of the tools, the definition of the needs and outcomes, and the testing of many models to find the combination that works best. Iterate.

7. **Maintenance is, still, the most important part.** Although it has been said many times in this report, as one of the practitioners we interviewed said – make sure they understand that launching is nothing compared to ensuring it works well, it continues to learn, and it does the job intended. What he said.

There are, I am sure, many more lessons that are either critical or that have been learned the hard way. While we may have had AI around for close to 40 years, we are still learning what it can (and cannot) and what it should do. We have been working with KM since the 1960s and are just starting to figure out a more efficient model for using it in the enterprise. AI is not far behind, finally.

Patience, however, is the most important lesson learned.



# Decision Framework for AI Technologies

An almost perpetual problem with AI technologies is deciding which one, or which ones, to embrace as an organization. I am constantly being asked about one or another – or worse, being asked to compare solutions from vendors that are quite different from each other. Comparing a machine learning solution with an NLP engine is not fair to either.

There are many aspects to making this decision one way or the other, but I found five elements that play into choosing one solution over another. The following is a decision framework – it allows you to make the decision but does not name any one solution. To be fair, no technology or tool can do what others cannot – they are all very capable, if you are aware of what you are trying to do.

## Outcome Sought

There are five customer-centric outcomes you should be seeking: understanding your customers, resolving their issues, generating actionable insights from interactions, managing expectations (both ways), and generating engagement (ultimate outcome). Each of these outcomes is the result of different actions; some of them are not traditional customer service actions but are becoming more rooted in service due to the actions customer service is adopting from other parts of the organization.

Each outcome will require a different approach to using and leveraging AI and KM and that will be determined by how intertwined information and the specific outcome are, and what are the informational needs of the processes.

If you are trying to solve an issue with, for example, understanding your customers better a tool like NLP may be better aligned as it does precisely that – breaks down the language into manageable components. If, however, you are more concerned with resolution, a classifier and automation tool may be better suited for that. AI and ML tools are all very different – not only in what they do, but how they do it – and they seek different outcomes.

A system that ensures the right information is communicated early to the customer – for the purpose of proactive communications and preventing an issue from arising – is not likely to deliver a lot of value in managing expectations (OK, maybe a little) but will deliver a significant amount of value in long-term engagement.

**Action Item:** *correlate the pain point or problem you are solving to one of the five outcomes above.*



## Expertise Required

An almost perpetual problem with AI technologies is deciding which one, or which ones, to embrace as an organization. I am constantly being asked about one or another – or worse, being asked to compare solutions from vendors that are quite different from each other. Comparing a machine learning solution with an NLP engine is not fair to either.

Each of the actions that customer service undertakes require different levels of expertise, different levels of information, and most likely different actions. The answer to the problem or question, or the resolution workflow that AI will enhance, is rooted by expertise and wisdom. Knowing where that expertise and wisdom resides, how to access it, how to leverage it, and how to maintain (optimize) it over time are the questions you should be asking next.

While an answer may reside in a knowledge base and be (mostly) static, at times it will reside with a subject matter expert (SME) and require a different path to acquire the information. Knowing which is which will give you the second clue as to which tools or technologies will work best

However, the concept of proper maintenance and proper allocation of resources to both creating as well as maintaining and extending the knowledge of the customer is critical to properly positioning AI in a KM setup.

Knowing what information is worth capturing, what is worth retaining, and how to use it and process it is what's going to make a difference. Until now, most KM systems simply stored all information and knowledge, then indexed it with the hope that it would become useful. That was never the case. Thus, we end up with untold numbers of unused "knowledge articles" that simply clutter the results and slow down the finding.

Understanding what knowledge is (distinguishing between content, knowledge, and data) and how it adds value to the overall process is

the first step. Finding how knowledge helps both the company and the customer reach the outcome is the goal.

**Action Item:** *determine who has the information you need, how you can find it, extract it, use it, maintain it and improve it. Use this information to plan your strategies.*

## Deployment Time

Many things go into understanding and making decisions on deployment time but the two most commonly used and that you should understand are training time and maintenance time.

If it takes months of training, as is the case in NLP or neural networks components, to reach an entry-level baseline status, and then more months to get to a better place, the tool may not be what you're looking for. On the other hand, if it requires constant supervision for maintenance yet you are short on resources, that may not be the tool that matters.

Each tool will require different levels of preparation and maintenance and you should understand these – deploying "within weeks out of the box" is not a metric, it's a marketing statement.

Understand what's behind the few weeks and how it works on an ongoing basis. These are the questions you will find provide the answer to which tools are best suited for a specific outcome you seek.

**Action Item:** *Plan for deployment, maintenance and iteration time for each specific tool when looking at which tools or solution from AI to leverage to enhance KM. The necessary investment may change your strategy given needed and available resources.*



## Enterprise Adoption

Culturally speaking, most AI tools will die from lack of adoption. While the technology may be there, and all the other elements may align, if the people in the organization don't use it or promote its use, it won't live for a long time.

These tools all require extensive follow-through and maintenance in addition to support and understanding how to achieve their objectives. None of these tools will be there over the long run if the ability to maintain them is not there, and the maintenance won't happen if the corporation does not adopt it and support it with budgets, time, and other resources.

**Action Item:** Ensure that whatever you are trying to do can be supported within your organization to determine if it makes sense to deploy.

## Cost

The “hidden” project-killer. While most of the tools available as services today are relatively cheap to implement, most of the costs of using them don't come from the tool itself but from the organizational costs of implementing and adopting.

AI is all about clean and useful data, and if that is not available in your organization, the path to get there is complex and expensive. Further, maintenance can cost as much as deployment – more in some cases – and require resources (people, budgets) you are not allocating today. As an example, traditionally KM tools require maintenance that increases in cost by 8% every year – is that cost something you can afford? How about double that (possible with more advanced tools that require more maintenance)?

It also requires data people, knowledge administrators, and potentially other people who are not in the organization today – can you get them?

**Action Item:** ensure you can get from your providers all the costs associated with deploying and maintaining a solution you are considering. *Licensing and deployment costs are just the tip of the iceberg, especially in agile and dynamic environments where maintenance and updates (of both information and processes) will make or break the project.*

These are the five areas you should understand, and they all combine in different ways for different tools – and different enterprises.

There is no “one size fits all” model for decision making; each tool and technology you will consider will have different answers to these questions.

It is almost impossible to have a chart or table that will guide you, and a matrix to highlight all tools and all these questions would be unwieldy – but knowing what questions to ask, and where to focus the strategy building, will go a long way to making a decision that will benefit your company.



## Case Study - BMW

### BMW Transforms Customer Relationships by Augmented Knowledge

BMW was a newcomer to the world of knowledge management in the customer space. Driven by the desire of customers for easy self-help and inconsistencies across touchpoints, they concluded that increasing their focus and investment in KM could help them fulfill their customers' expectations better. They identified key customer touchpoints (external, partner and internal; digital and physical) and touchpoint support functions as well as internal functions where knowledge would deliver value – and then they set out to look for a solution.

Right on the heels of the explosion of AI offerings in the marketplace, they felt as if many of the solutions were too complex, or too demanding of resources, to be of value. Training a system for months on end, only to find the maintenance requirements equally taxing on an ongoing basis was not the solution they needed.

Through exploration they found myriad knowledge sources that would need to be aggregated, different departments and stakeholders that needed to be acknowledged and included in the planning, and a knowledge generation transformation that needed initiating. They looked at solutions available in the marketplace and figured they needed a system capable of augmenting knowledge, not just managing it.

According to Suzanne Gray, "we knew we could operate far more effectively (reducing repeat contacts across touchpoints or escalations) by focusing on customer needs wherever they were (including search engines). We wanted our effort to be on the answer, rather than on the systems to deliver it". Their solution was to use a solution with AI to simplify finding the right content, increase content relevance, and refine the content via easy maintenance.

Doing a piecemeal deployment, touchpoint by touchpoint and brand by brand, they found the kinks as they went along and learned every step of the way. They also found that the iterative process was invaluable to see where new layers of value were being created by the AI capabilities of their KM solution.

We knew we could operate far more effectively by focusing on customer needs. **We wanted our effort to be on the answer rather than on the systems to deliver it.**

**SUZANNE GRAY**

BMW

One of the most valuable lessons they learned was that users were never just asking for a piece of information, but that context (and intent) were the real proposition to deliver value. By focusing their efforts around understanding their customers better – including the real reason why they were asking the questions – and by using those insights to augment the delivery of knowledge to customers by also catering for their next most likely need they could focus on delivering the most accurate answer and improve customer journeys.

By monitoring behavior as well as feedback, they could keep improving delivery of the right knowledge, at the right time, in the right place to different stakeholders and learn along the way which piece of knowledge was truly helpful and valuable and which did not work as intended. This is a part of transforming the automotive industry mindset around the relationship with customers in a digital, data-driven age, where the traditional responsibility for serving the customers' needs has sat with franchised dealers.



## Next Steps

### Where to start?

That's the perpetual question I get at this point. Even the practitioners in this study, most of whom have already implemented AI and KM, were interested in knowing if I had a methodology to make it easier, faster, better to continue the iterations.

This report summarizes the two things you must know: what is AI (as it should be used in organizations), and how it benefits and augments KM for those doing it. It removes a lot of the hype and uncertainty of understanding what the outcomes should be, what the results will take to accomplish, and what the model looks like.

The rest, the application, is up to you. Your situation is unique, and your setup is unique. No one can tell you whether it will or won't work for you – you must commit to trying it.

Start slowly, test and pilot, learn by trial and error on small projects. That's always the best way to learn. Use the lessons learned from implementing other disciplines (most notably KM) prior and leverage all the resources you can get your hands on.



THINKJAR  
estebankolsky.com

Author:  
Esteban Kolsky,  
ThinkJar

Use case studies, lessons learned, methodologies, and frameworks as those presented in this report to create your strategy. Create a strategy and modify it as you go along.

More importantly, have patience and always, always know what you are trying to do.

Anything else, just ask.

Predictive models for KM did not make sense to us until we figured out the secret: it's not about repeating the interaction verbatim, but about having a **"close enough" interaction** that could use the same content.

**KM MANAGER**  
Global Commercial

### Europe, Middle East & Africa

info.emea@verint.com  
+44(0) 1932 839500

### Americas

info.apac@verint.com  
1-800-4VERINT

### Asia Pacific

info.apac@verint.com  
+(852) 2797 5678

# VERINT.

The Customer Engagement Company™

Unauthorized use, duplication, or modification of this document in whole or in part without the written consent of Verint Systems Inc. is strictly prohibited. By providing this document, Verint Systems Inc. is not making any representations regarding the correctness or completeness of its contents and reserves the right to alter this document at any time without notice. Features listed in this document are subject to change. Not all functionality is available in all configurations. Please contact Verint for current product features and specifications. All marks referenced herein with the ® or TM symbol are registered trademarks or trademarks of Verint Systems Inc. or its subsidiaries. All rights reserved. All other marks are trademarks of their respective owners. © 2019 Verint Systems Inc. All Rights Reserved Worldwide. 5.2020