

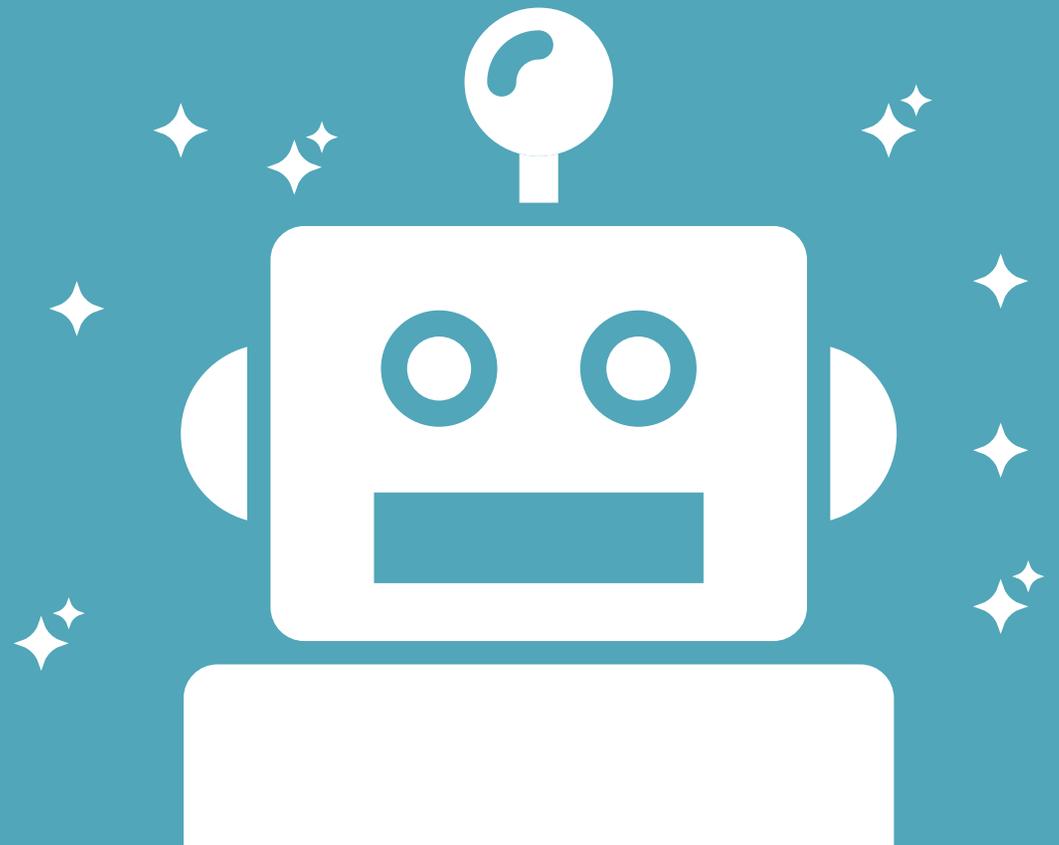


# Bullshit, Hype, and a Little Magic

*How to Make Sense of It All When Buying A.I.*

## *Products*

By Rob May, CEO and  
Co-Founder of Talla



# TABLE OF CONTENTS

Introduction.....	3
Should You Even Consider Buying A.I.?.....	4
Cutting Through the Bullshit - What Really Works?.....	5
• Sorting the “real” A.I. from the fakes.....	7
Parsing the Hype - How to Ask the Right Questions.....	9
• Ask about training data.....	9
• Ask about ramp-up time.....	10
• Ask about error-correction.....	10
Finding the Magic - How to Spot the Real Potential.....	11
• Most promising A.I. domains today.....	12
Summary.....	13
APPENDIX: The PAC Framework.....	14

## INTRODUCTION

I've been involved in selling enterprise software in some way, shape, or form for the last 12 years. As the CEO of an artificial intelligence software company, angel investor in over 30 other A.I. startups, and author of the popular [InsideAI newsletter](#), I've had an inside look at the A.I. sales process across many companies and my conclusion is that the sales process for A.I. software is a confusing, hype-filled roller coaster for most buyers.

The companies developing many of these A.I. products make wild claims that they can't back up, or oversell decades-old data science techniques as "A.I." just to seem sexy. But despite the amount of snake oil in the industry, buyers also know that artificial intelligence is making all sorts of extraordinarily powerful new solutions available - they just lack the expertise to distinguish the real magic from the snake oil.

The goal of this eBook is to lay out the most important factors to consider when buying an A.I. product. What questions should you ask? How is A.I. different from the software you bought before? How do you know what works? And most importantly, is the return on investment worth it?

Bottom line: How do you cut through the hype and the bullshit to find the products that actually deliver on some A.I. magic? Find out below.

(And if you're impatient, just skip to the appendix. All the hype-free, no-bullshit magic is there.)

## SHOULD YOU EVEN CONSIDER BUYING A.I.?

Yes. Yes. Yes.

Admittedly, I am biased here, but let me explain why you need to dip your toe in the A.I. water now. There are three key reasons.

### 1. Some A.I. software really can do magic

A.I. can have real, significant benefits for certain use cases, as we will see in later sections.

### 2. You need to build A.I. expertise

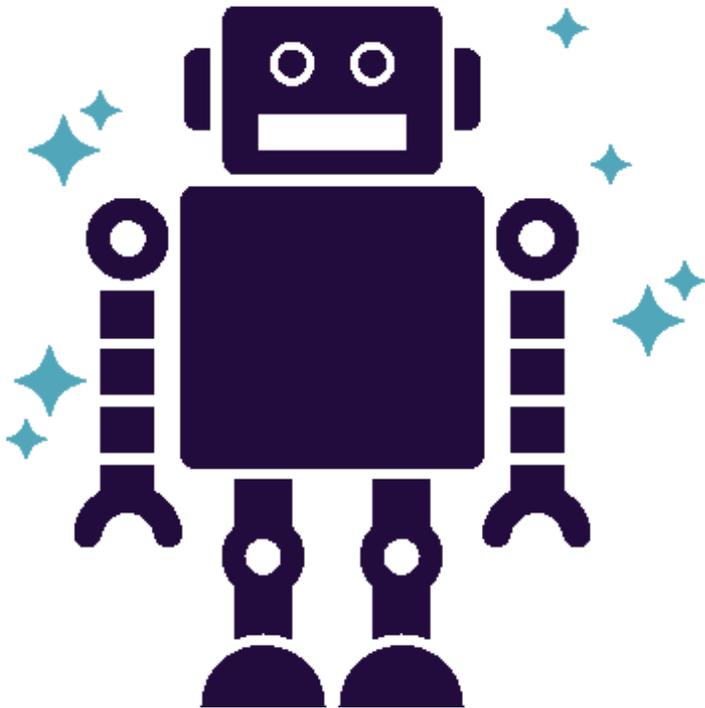
A.I. is a complicated space that is changing rapidly. The more you start to build internal expertise with A.I. products, the better prepared you will be to navigate the coming deluge of A.I. software.

### 3. You have to start your A.I. flywheel

Most A.I. systems are like flywheels; the longer you run them, the more effective they become. A.I. models get smarter with more data over time. The sooner you get started, the sooner your data models get “smart” and the harder it will be for your competitors to catch up.

I can't stress this last point enough, so let me say it again. **Start looking into A.I. today, or you will be left behind.** You cannot deploy an intelligent system one year after your competitors and catch up the way you can with regular software. Time matters for artificial intelligence software much more than it normally does when deploying traditional software.

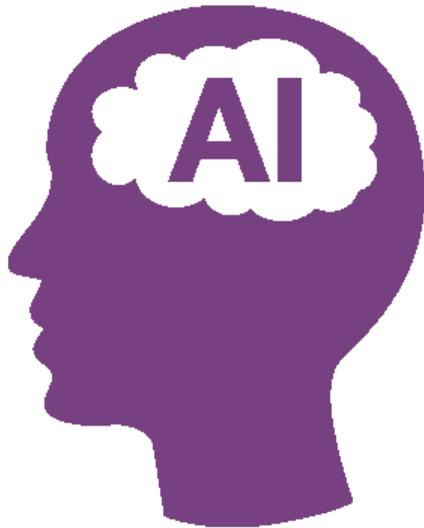
If you want to know where to start, check out the PAC Framework in the appendix. PAC stands for *Predict*, *Automate*, and *Classify*, and is a simple way to think about where A.I. can impact your business the most.



## CUTTING THROUGH THE BULLSHIT - WHAT REALLY WORKS?

If you have decided you need A.I., that is a great decision, but now you must navigate a landscape of vendors that includes:

- Tools pitched as A.I. that aren't really A.I.
- Tools where the A.I. isn't obvious
- Tools that don't really work
- Tools that use humans to do tasks in hopes someday they can get enough data to let A.I. do the work
- A small group of tools that are actually magical



The obvious question that arises is - what really works?

*Gödel, Escher, Bach* author Douglas Hofstadter once remarked that “artificial intelligence is anything that has not been done yet.” With that, he meant that the term “A.I.” is often misapplied to any cutting-edge computing technology, and that any A.I. technique that becomes widely productized is no longer considered “artificial intelligence” - because what’s ordinary cannot ever be “real” A.I.

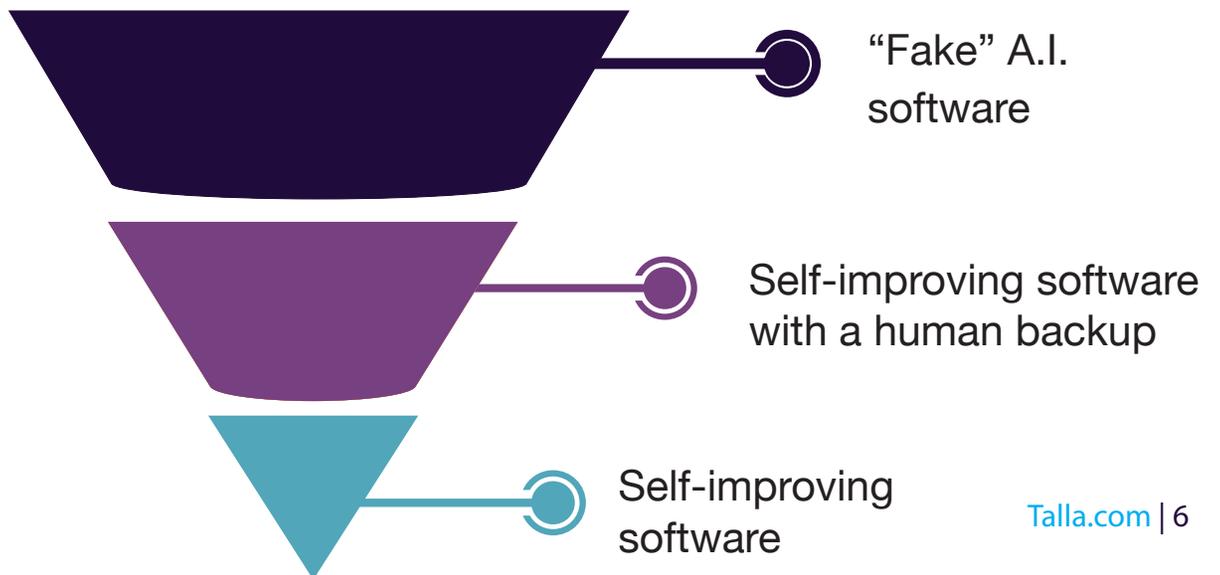
Arthur C. Clarke famously said that “any sufficiently advanced technology is indistinguishable from magic.” Artificial intelligence is the term many laypersons (and far too many marketers) use to describe any software that seems “magical” - regardless of whether any actual A.I. is under the hood. Moreover, many widely used software solutions and features - from Facebook’s Edgerank news feed algorithm to Apple’s Siri virtual assistant - use A.I., but aren’t typically marketed as artificial intelligence software.

The way I think about A.I. is that a system is “intelligent” if it can change its own programming or responses over time, based on experience with the world.

A looser definition of A.I. would be the same as the previous sentence, but to include humans as part of the process. For example, a tool that schedules meetings via email, that makes many decisions on its own but defaults to humans when it can’t understand what it’s being asked to do, would still be considered A.I. under the second form of the definition, but not the first.

There is also a third version of A.I., which is *over-marketed artificial intelligence*. This definition crops up when a company is using a basic data science technique like support vector machines that is well known and understood - and may have been in their product for years - but is now marketed as A.I.

When you encounter “artificial intelligence” software, just be aware that the term could mean any of these three definitions, and your evaluation of that software should begin by determining which of the three A.I. definitions the vendor is using.



When you encounter “artificial intelligence” software, be aware that the term could mean any of these three definitions.

## SORTING THE “REAL” A.I. FROM THE FAKES

Whether an artificial intelligence solution wants or needs a human backstop is often a matter of preference, or what’s appropriate for a particular use case. The real trick is eliminating software that doesn’t use actual A.I. because using that software won’t build your A.I. expertise or start your A.I. flywheel. You can often sniff out fake A.I. by asking what type of artificial intelligence technology a software solution employs. Solutions that don’t use any real A.I. typically can’t or won’t answer the question.

Here is a table of the commonly used types of A.I.:

Artificial Neural Networks	Logistic Regression
Deep Learning	Probabilistic Logic Networks (PLNs)
Evolutionary Algorithms	Support Vector Machines (SVMs)
Hierarchical Temporal Memory (HTM)	Word Vectors/Word Embedding



The real trick is eliminating software that doesn’t use actual A.I.

These technologies can all be applied to almost any task or industry, but there are certain areas where A.I. is more useful or more directly applicable. If you want to be sure you're buying a real A.I. solution, it's often useful to ask what the "A.I." in the software is being used for. Here are some common A.I. technology applications given the current state of the art.



If you're looking to build your A.I. expertise and start your A.I. flywheel, you need to buy "real" software.

Automation	Natural Language Generation (NLG)
Machine Vision	Predictive Analytics
Natural Language Processing (NLP)	Speech Recognition

While there are plenty of "magical" solutions that don't strictly employ A.I., if you're looking to build A.I. expertise and start your A.I. flywheel, you need to buy "real" artificial intelligence software. These questions will help you do that.

The more difficult task is in finding the *right* A.I. solution that can actually deliver magical results for your business.

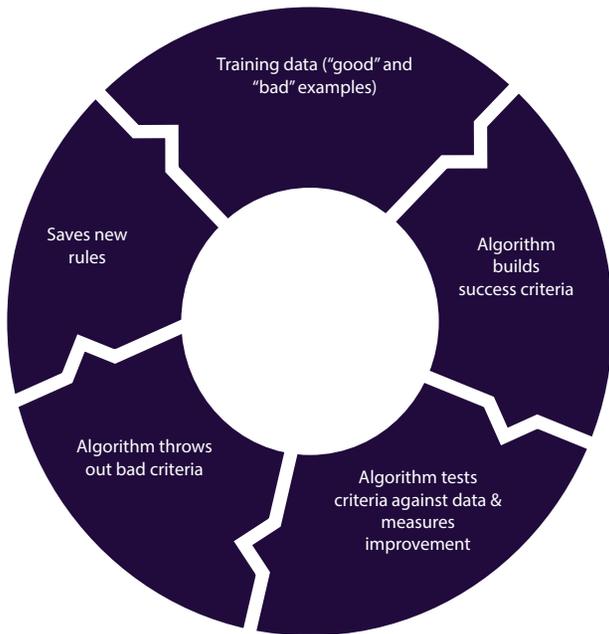
## PARSING THE HYPE - HOW TO ASK THE RIGHT QUESTIONS

In general, there are three types of questions you need to ask to find a truly valuable artificial intelligence solution. These questions help you identify the underlying type of A.I. used, how long that A.I. takes to get up to speed and become useful, whether the A.I. improves over time, whether the improvement process requires humans, how the A.I. gets fixed if it makes a mistake, and how the A.I. performs against a non-A.I. version of similar software.

### Ask about training data

The place to start asking questions for an artificial intelligence product is the data. A good opening question is something like, “*what data did you use to train your models?*” Much of the newest A.I. software uses neural networks, or deep learning (a specific type of neural network), and these models need training data. If the software performs a relatively generic task - like [machine translation](#) - the models may be pre-trained. If the task is specific to your data, you want to understand whether the tool is useful “out of the box,” or only after it has been trained on some data. Ask questions about what data is used to train the A.I., how the A.I. learns, and whether A.I. adapts to new data.

If an artificial intelligence solution requires training on *your* data, will the vendor need a data dump to get started, or can the tool learn as it goes? One big misconception about A.I. products is that they are all smart out of the gate - they may not be. It depends on what the A.I. needs to learn, and where the data to learn it comes from. Some A.I. solutions may take more time than others to show value.



If your own data is required to build training models for the A.I. solution, you want to make sure you ask questions about data rights and privacy. Are the A.I. logic models unique per customer, or trained across all customers with aggregated data? If the training data is aggregated, how is it anonymized before being added to the overall training pool? If you leave the vendor, do you still own the model you built, or does it belong to the vendor - with the opportunity to resell it to your competitors?



### Ask about ramp-up time

The second set of questions to ask your A.I. vendor deal with learning and performance. How long does it take for the tool to perform at an acceptable level? Does the A.I. require you to use it for a few days or weeks first? If the software learns more over time, does it use a model that does so automatically, or does it require feedback from humans?

### Ask about error-correction

The third area to ask questions about is error-handling. Intelligent systems often have probabilistic outputs, meaning they give an output more like a weather forecast (which is occasionally wrong) instead of a clear logical output like most software. For example, a weatherman tells you there's an 80% chance of rain, and an A.I. tool may tell you there's an 80% chance an image has been photoshopped and isn't real. You need to understand the error rates of the A.I. models used in your software - how often their forecasts or predictions are wrong - and what should happen when you think it is incorrect. How do you grade the performance of an A.I. solution and guide it to improve?

## FINDING THE MAGIC - HOW TO SPOT THE REAL POTENTIAL

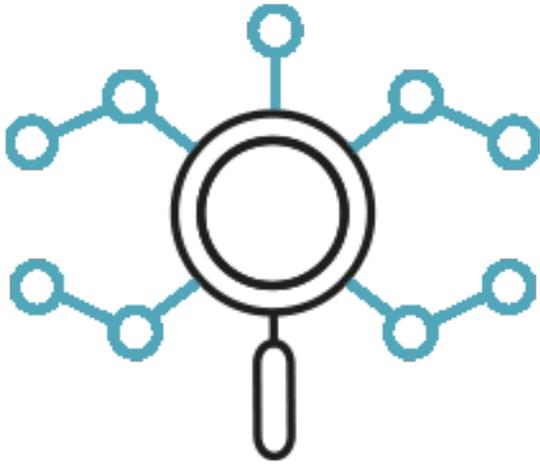
The best applications of artificial intelligence at the moment are problems where the solution has a lot of rules, but where those rules would take too long to write down, and would be complicated because they contain lots of exceptions. If you have a large data set, you can feed it to a neural network model and the model will figure out the rules on its own. That's the beauty of A.I.

Beware of rule based systems without A.I. - they tend to be "brittle," meaning that if you don't closely follow their interaction or input expectations, they can easily fail.

Here are some general rules for spotting potential in different applications:

**Chatbots** - Are they scripted, or do they actually parse language? You can tell by asking the same question two different ways. "Who provides our health insurance?" and "Who is our health benefits provider?" should return the same answer for an HR chatbot. If the results are scripted, instead of real natural language processing, they may not. Lean towards real NLP bots - they will perform better in the long term.

**Predictive Analytics** - These tools are most useful in sales and marketing. Test to see if they generate insights better than your sales reps, and if those insights prove to be accurate and valuable. A good test might be to ask your sales reps to estimate the probability of close within a certain time frame of 10 deals they are working. Then get the same probability from the A.I. tool and see which is better.



**Classification Tools** - Classification tools should be able to work well out of the box because they can be trained on data from other customers, while still keeping that data private. If your tools are meant to divide something into groups or classify them, test to see how quickly it can learn a new category or pattern by providing it with something new. For example, for a tool that classifies financial transactions automatically, try to give it a complex transaction.

**AI Infrastructure** - There are many infrastructure tools available that aren't intelligent themselves, but they enable smart tool building. When evaluating these tools, make sure that they're flexible, support all major neural network use cases, and that they're run by either well-funded private companies or crowdfunded communities, so they're around for a long time.

Artificial intelligence tools are very early as a market, but some of them can show you small bits of magic, and those show where this industry will go over the coming years.

## MOST PROMISING A.I. DOMAINS TODAY

**Machine vision** - Machine vision isn't just for robots; building software that can parse and recognize image or video data has an extraordinary number of applications. Using artificial intelligence to read text, recognize faces, or identify objects in videos or static images has proven immensely useful, and the technology and techniques behind these solutions is among the most rapidly maturing domains of modern A.I. software.

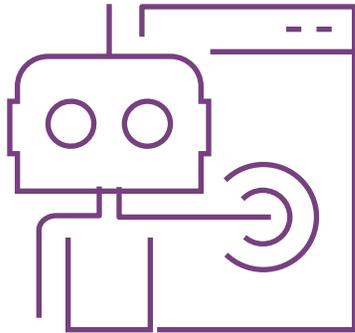
**Chatbots** - Developing A.I. tools that can successfully converse with humans - recognize written or spoken requests or instructions, and respond in a conversational written or spoken form - is an extremely rich area of A.I. development. Speech recognition and natural language



processing now come standard with every smartphone, and the A.I. behind Siri, Alexa, and Cortana is now making its way into broader enterprise applications.

**Predictive Analytics** - Using artificial intelligence to make more accurate and more specific predictions is one of the oldest, yet still the most rapidly improving, areas of A.I. development. While financial firms have used predictive analytics to improve their investment and stock trading systems for decades, these algorithms are now being applied and improved for application in medicine, travel planning, and even sports analytics.

If you want to invest in an artificial intelligence solution that likely contains some real, useful A.I. with highest potential ROI, focus on solutions within one or more of these domains.



## SUMMARY

Many solutions marketed as “artificial intelligence” are bullshit, and many more are greatly overhyped, but some A.I. tools really do deliver some magical results. By asking about the underlying technology, getting clear answers on how and how long it takes to implement an A.I. solution, and investing in A.I. that is put to a proven, well-understood use case, you can help ensure that any A.I. solution you buy will have the best chance for a real return on investment.

Above all, by investing in the right A.I. solution now, you’ll not only get a good A.I. ROI in the present, but you’ll develop the A.I. expertise - and start the A.I. flywheel - that will set you up for greater success in the future.

## APPENDIX: THE PAC FRAMEWORK

If you want to apply A.I. at your company, where do you start? The PAC framework is a good tool to use to identify the most promising areas. PAC stands for: *Predict*, *Automate*, and *Classify*. While A.I. can do many things, these three comprise the bulk of use cases today.

To get started with the PAC framework, make a chart with three boxes on the vertical axis for three areas of your business: *Customers*, *Product*, and *Operations*, and three horizontal boxes for: *Predict*, *Automate*, and *Classify*. In each box, list something that fits both dimensions. For example, at the intersection of “customers” and “predict,” what might be a good use case for predicting something about your customers? Maybe predicting attrition, or repeat purchases? Below is an example.

	Customers	Product	Operations
Automate	<ul style="list-style-type: none"><li>• Lead generation</li><li>• Sales prospecting</li><li>• Call follow up</li></ul>	<ul style="list-style-type: none"><li>• Onboarding and training</li><li>• Bug resolution workflow</li></ul>	<ul style="list-style-type: none"><li>• Common monotonous workflows</li></ul>
Classify	<ul style="list-style-type: none"><li>• Which customers are most profitable?</li></ul>	<ul style="list-style-type: none"><li>• Customer input</li><li>• Bug classification</li></ul>	<ul style="list-style-type: none"><li>• Information and expertise</li></ul>
Predict	<ul style="list-style-type: none"><li>• Which deals will close?</li><li>• Which customers will churn?</li></ul>	<ul style="list-style-type: none"><li>• What does your customer want to do next?</li></ul>	<ul style="list-style-type: none"><li>• Shortfalls</li><li>• Employee attrition</li></ul>

Once you have your framework filled in, ask yourself two questions about each item in each box.

1. How valuable is it to the business if we did this?
2. Do we have a large data set that we could use to train an A.I. for this?

Start with the one or two things that rank highest across both questions. That will give you a good starting point to learn and apply A.I. in a useful situation. From there, you can expand as you understand A.I. tools and techniques much better.

## ABOUT TALLA



Talla's AI-powered knowledge base and intelligent agents bring IT and other business teams into the future of automation and insights. The Talla platform uses machine learning to surface relevant information in the right context, and keeps regularly changing information up-to-date. Use Talla to onboard new employees, for policy and procedure management, and more. Deliver and collect information through chat in the TallaChat application or through Slack or Microsoft Teams. Get started today at [Talla.com](https://Talla.com) or by calling 617-517-4156.