monetate

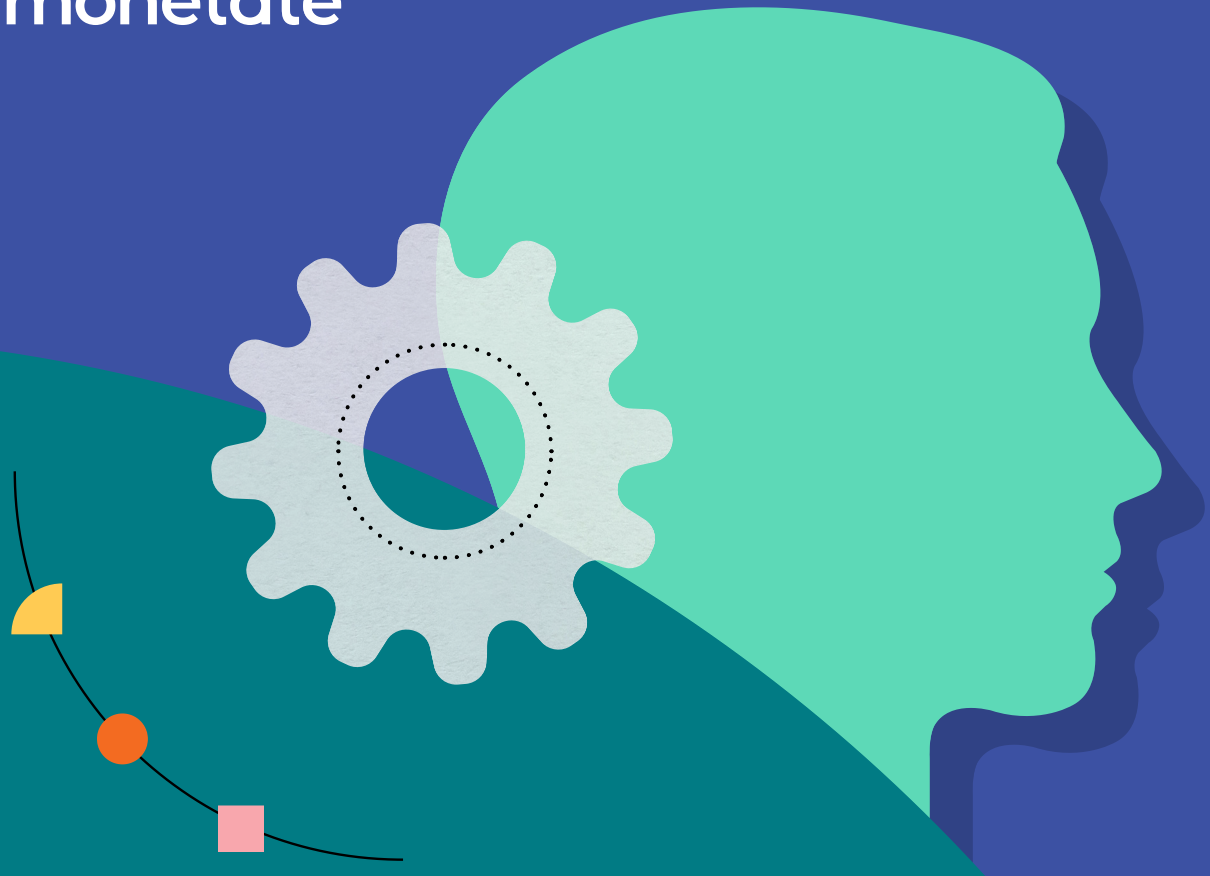

Evaluating Your Next ML Personalization Engine:

Don’t Be Fooled By Deep-Learning

Table Of Contents

Introduction

Marketers are accelerating their investment in digital transformation. After a year of accelerated online activity, online channels are becoming a major focus. To make a digital strategy scalable and cost effective, companies need to combine a variety of different components across data and technology to deliver the best possible customer experience. The glue that binds these components is personalization at scale – the ability to deliver personalized experiences to many people across channels.

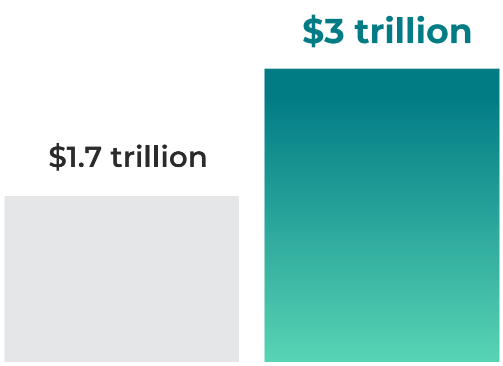

Personalization at scale has the potential to create $1.7 trillion to $3 trillion in new value in the next few years, with retail making up the largest portion of that opportunity.

McKinsey Global Institute

McKinsey Global Institute projects that personalization at scale has the potential to create $1.7 trillion to $3 trillion in new value in the next few years, with retail making up the largest portion of that opportunity.

At the heart of personalization at scale is machine learning (ML), a combination of data and technology that can make rapid decisions based on many insights. ML has the ability to deliver personalization that uses real-time insights to adjust the customer experience, and improve over time to become more valuable with each interaction.

ML as a technology is beginning to mature with marketers, who are now hearing about variations of ML like "deep learning." Deep learning, (also known as neural learning) refers to a complex form of machine learning that imitates the functions in the human brain. It sounds exciting, but it isn't necessarily the best path for marketers who are looking for ROI and accessibility. Rather, well targeted, more approachable ML can deliver dramatic improvements to customer experience and profits.

In this paper, we'll help you to understand the different types of ML, why they're important, the key questions you should be asking, and the answers you should hear when speaking to a vendor offering ML-powered personalization capabilities.

Key Takeaways

MACHINE-LEARNING (ML) IS KEY TO KEEPING & WINNING CUSTOMERS

It's no secret that Machine Learning is essential to deliver engaging, personalized experiences at scale. With ML, even a small team can scale, pivot, and be truly omni-channel to serve increasingly discerning customers.

THERE WILL BE A TIME FOR DEEP-LEARNING, BUT WE'RE NOT THERE YET

Deep learning has it's uses but for modern-day marketing applications, the additional complexity of deep learning is not only unnecessary, but it's also not very practical.

ROI AND ACCESSIBILITY ARE A MARKETERS' NORTH STAR

A more approachable (easy-to-use and interpret) ML can deliver dramatic improvements to customer experience and profits. Deep Learning sounds exciting but as marketers, we should prioritize ROI.

IT'S ALL ABOUT EXPLOITATION VS EXPLORATION

A good ML engine, whether it's deep or reinforcement learning, should help you quicklyscale the process of exploitation and exploration, alongside allowing you to understand the results - Balancing learning with performance.

Definitions

1-TO-1 PERSONALIZATION

1-to-1 Personalization describes the practice of delivering a unique, optimal digital experience to each customer using all available data from 1st- and 3rd-party sources. To deliver a customized experience to every visitor across channels in real time, you need a 1-to-1 personalization approach which requires rapid data aggregation and analysis, cross-channel deployment, and ML. 1-to-1 personalization is the output of ML-powered personalization done well.

ARTIFICIAL INTELLIGENCE (AI)

Alan Turing defines AI as: ".. the science and engineering of making intelligent machines, especially intelligent computer programs."

AI helps marketers understand vast amounts of customer, business, and market data, better anticipate customer intent, make smarter, data-driven decisions, and ultimately improve the customer experience.

DATA

Data refers to insights and information about a variety of factors relevant to marketing. Customer data, which is often stored in a CRM system or customer data platform, includes personal information such as email address, loyalty program status, and more. Business data can include product inventory, details about margins, fulfillment routes, and more. Market data can include a host of external insights from weather to economic trends that are relevant to a business and its customers.

For personalization and targeting, customer data is the biggest focus. Customer data that's collected directly by the marketer and stored in their own databases is called "first-party data" while data that is bought for use by other companies is called "third-party data." Each type of data is subject to a variety of privacy requirements to ensure it is collected, stored, and used in a way that does not compromise the safety of the individual. Marketers can also leverage in-session or behavioral data from within a personalization solution, this can include a huge array of insights from purchase history, brand affinity, device type, geography, search behavior, etc.

CONTEXT

Context is the technical phrase for *what we know about the visitor — it's the data used to make decisions.* Things like time of visit, device, past purchase, viewing behavior, customer segment, any third-party data point, etc. are all context variables that can be fed into a decisioning ML engine to learn and make decisions.

MACHINE LEARNING (ML)

ML is a subset, or branch of AI. It's the process of a tech system learning from data inputs and using those inputs to recognize patterns, develop conclusions, and in some cases, produce outputs. The basic flow of ML technology is: one, using collected data to generate an algorithm, two, gathering feedback, and three, applying learnings in the form of an output or action.

ML can offer a "point solution" — a single algorithm that makes specific decisions within a single page on a website, for example. It can also be very complex, using models to combine insights across a variety of touchpoints for a highly orchestrated outcome — for example, enabling driverless cars.

The value of ML is in automating the decision-making that can come out of AI analysis and enhancing marketers' ability to use large amounts of data within marketing processes.

The key differences between AI vs ML are:

1. A ML algorithm will only ever understand what it was trained to do, and therefore doesn't understand situations outside of its training. It can recognize patterns within the realm of existing training (like spotting fallen trees on a road) but will not understand and learn new hazards unless a human re-trains it.

2. The goal of ML is to understand the structure of the data and allow us as humans to understand it, as well. For this reason, ML forms the basis of most AI systems and is most widely applied specifically in marketing use cases, where understanding and insights are key.

NEURAL NETWORKS

Neural Networks mimic the activity of a brain through a set of algorithms. They add a clustering and classification layer on top of the data you already have. All ML technologies use a neural network, the differences are around the 'depth' or the number of layers in a neural network.

Why Machine Learning?

The research firm IDC notes that the pandemic has caused marketers to squeeze years of digital transformation into a shortened timeframe. Marketers are making investments to “shore up marketing technology, integrate data and analytics, and add digital skill sets.” The goal is to keep up with digitally savvy customers. Mobile search and shopping, social shopping, same-day shipping, streaming sales — all of these consumer behaviors require new capabilities for companies to keep up.

ML is an essential element — helping companies ingest and process substantial amounts of data quickly, test and learn in real time, and automate complex processes. With ML, even a small team can scale, pivot, and be truly omni-channel to serve their demanding customers.

The benefits of ML for marketing teams are many. Here are a few of the most important:

SCALING UP

Before ML, “personalization” would mean delivering targeted messaging, or product recommendations to different customer groups. This is what many now refer to as “segmentation”. Segmentation is a valuable tool for marketers but is limited in its flexibility to deliver relevant experiences to every customer.

Separate experiences for your top customer segments may be lucrative and worth the effort, but the creation of experiences and the rewards from slicing and dicing ever smaller groups of customers diminishes as you create smaller and smaller segments. Trying to deliver rules-based segmentation to all your visitors on just one channel is nearly impossible, never mind trying to cover cross-channel.

With a good ML decisioning engine, you can achieve individualized personalization at scale, and use a wider amount of data to develop new segmentation that can easily adapt to changing behaviors. ML automates the process and speeds it up, allowing marketers to drop in different experiences (content /messaging/ different product recommendation models, etc), set a goal (e.g., a conversion threshold), and let ML do all the heavy lifting.

GETTING MORE FROM DATA

Most customer-facing companies have a wealth of customer data that is sitting there waiting to be utilized, from sales and loyalty data to shipping, returns and browsing data. ML makes it feasible to feed this valuable data into a decisioning engine to help influence decisions. A good ML engine will give marketers out-of-the-box targets to utilize from day one, but also provide flexibility and data-openness to allow for new inputs.

Benefits of ML for marketing teams

- Scaling Up
- Getting More From Data
- Uncover Insights
- Deliver high performance

UNCOVER INSIGHTS

A marketer-friendly personalization engine will do more than just deliver experiences, it should also share insights around why those decisions have been made, what experience has resonated, and for which types of customers.

These types of insights will not just help with reporting on the performance of individual experiences, they can also aid in the understanding of your customers to be used across initiatives beyond just eCommerce.

DELIVER HIGH PERFORMANCE

ML-powered personalization maximizes performance. AB testing and segmentation experiences are important results-drivers but are not always the right fit for more complex needs, and they will never compete with the returns gained from 1-to-1 personalization.

When we say increased performance, this covers ROI from your personalization engine, but also from your content. Your content is valuable, a well-tuned personalization engine will ensure that each creative asset is displayed to the most suitable customer at the most optimal time.

ML-powered personalization maximizes performance by putting an experience in front of each visitor that is most likely to result in the goal metric you have decided whether that is reducing bounce rates or increasing conversion. Performance can be determined by specific goals that the algorithm can be trained to, from ROI to contextual relevance, efficiency, and more.

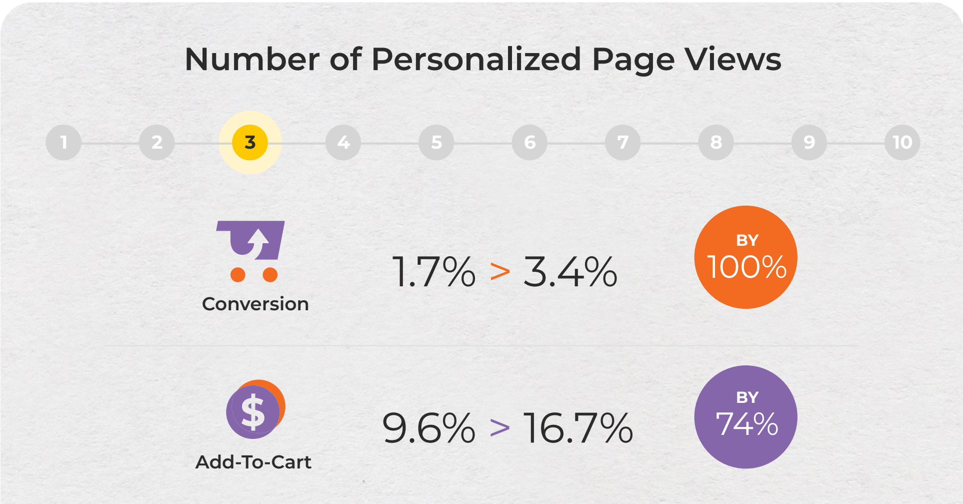

Source: Monetate Ecommerce Quarterly Report, 2019

Different Types of Machine Learning

There are three components to a well-working ML engine: data processing, model building, and deployment & monitoring. The middle section, the 'model', is the ML algorithm, of which there are three common categories:

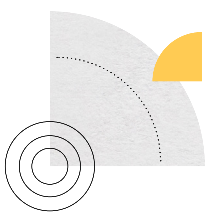

Classic Ml

Classic or traditional ML technologies do not use deep neural networks (unlike deep learning), nor do they use a reward scheme such as in reinforcement learning (see below). Classic ML has a more limited capacity to capture information about training data.

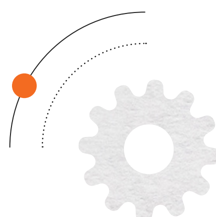

Reinforcement Learning

Is a subset of classic ML, and can be thought of as the next generation after classic ML. With Reinforcement ML, human nature is applied to the algorithm in a reward scheme (like a carrot). For example, setting a 'reward' as a goal metric like conversion helps the machine know to optimize performance against a set goal to 'reinforce' learning. The *reward* can be thought of as feedback and is used to teach the machine how to optimize its performance. Interaction is needed for this kind of learning to 'reinforce' learning. This subset of ML is well suited to marketing applications that must be goal-metric orientated (more on this later).

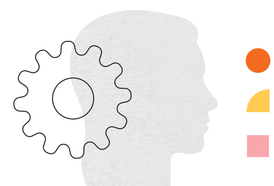

Deep Learning

Also, a subset of ML. deep learning uses neural networks with three layers or more to learn more complex relationships and make more accurate decisions. The 'deep' in deep learning refers to the depth of layers in a neural network. The number and depth of layers in a neural-network correlate to the level of complexity of relationships it can understand. Deeper networks can result in more complex decisioning, but also requires more data and therefore often more time to learn (more on this later).

The significant difference between Deep Learning and Reinforcement Learning is in how each algorithm learns and how much data each type of algorithm uses. The data element is key, as this is one of the key downsides to deep learning.

The Multi-Armed Bandit Problem

To understand the need for advanced ML methods, we first need to understand the multi-armed Bandit problem (MAB). Named after a hypothetical gambler at a row of slot machines. The MAB speaks to the problem or need for advanced ML technologies like reinforcement learning and deep learning as it **explains the *exploitation vs exploration* dilemma.**

MAB is named after a hypothetical gambler at a row of slot machines. It's important to set out and test each machine, noting the payout frequency and quantity, to understand which ones have the highest probability of winning. But this takes time and money, and testing comes with its own cost (direct costs and opportunity costs.)

When there is a constrained set of resources, a marketer needs to determine how to best allocate them to maximize the outcome. For marketers, the different "slot machines" represent different buttons, images, messaging, or recommendations algorithms. Instead of arms to pull the slot machine arms, there are data-driven algorithms. The goal of the test is to determine how to balance between exploitation (sticking to playing this current slot machine), vs. trying out additional slot machines to see if they will return a higher yield.

Good ML allows a marketer to exploit and at the same time explore whether there are more productive areas to exploit. The important thing to remember with ML is that **"you have to pay to learn."**

Solving For The Mab Problem

Data scientists have created several different solutions, or models to tackle the MAB. Some of the most common are:

Epsilon-Greedy: Always pulling the slot machine lever with the highest known payout, and choosing a random arm a fraction of the time

Decayed Epsilon Greedy: As you become more confident of which is the best arm to pull, you dedicate less resources to exploration

Upper Confidence Bound: The most commonly used method for solving the MAB, this method is based on the principle of optimism in the face of uncertainty which, means the more unsure we are about an arm, the more crucial it is to explore that arm (we assume that the unknown will reward us.)

Thompson Sampling: The Thompson Sampling (TS) method, also known as Bayesian, is gaining the most traction in recent years. As stated by method's namesake William R. Thompson, *"Thompson sampling is a method of solving the exploration-exploitation dilemma in a MAB problem." (exploring)* and rewarding *(exploiting)*, allowing *(in our case)* the marketer to identify which option to pursue in order to maximize returns. In scenarios where there is a clear winner, TS will favor the winner. This is one of the reasons that TS has gained traction as the learning base of choice for marketer focused AI-based learning technologies.

Probabilistic: *Unlike Upper Confidence Bound (where we're basing decisions on nonparametric, model-free assumptions), Thompson sampling allocates traffic according to the probability each variant is best.*

Accommodating delayed feedback: *Randomization naturally handles delayed feedback (minutes/hours between seeing an experience and conversion, for example) more gracefully than deterministic algorithms.*

Better empirical evidence: *Thompson sampling has been shown to be both theoretically optimal and preform very competitively on a variety of real-world data sets both in our five years of experience and the academic literature.*

'Explainability' of the traffic distribution: *Since traffic is allocated according to the probability that each variant is best, the results of the algorithm are easy to understand.*

For the machine to learn what experience will resonate best for each visitor, it must collect information on the performance of that experience (say number of clicks, or conversions), before the machine decides which experience is most likely to resonate best for each visitor. Therefore, the most robust, marketer focused ML powered personalization engines will solve for the 'problem' of having to pay to learn by using Thompson Sampling (AKA Bayesian).

Bayesian, or TS, allows the machine, and us as marketers to be efficient and quick to learn, therefore faster ROI, alongside having a view of what is going on under the hood. TS has gained so much popularity because it isn't a black box. You can (should) be able to report on why certain decisions have been made, allowing you to know where your creative, dev, or product teams should allocate resources.

Due to the nature of TS (matching the number of 'pulls' to the probability of the yield), the machine can make decisions earlier before we have a 100% answer. This methodology allows us to see ROI faster.

Deep Learning vs. Reinforcement Learning: What's Right For You?

Reinforcement learning is effective in most modern-day marketing applications, making the bells, whistles, and additional complexity of deep learning often unnecessary, but also not very practical. For this reason, the most advanced software engineers and data scientists use traditional models more extensively than deep learning, even with all the deep learning hype.

In most situations so called 'Deep Learning capabilities' will be classic or reinforcement learning models, with a heavy-handed marketer choosing ambitious product names. In the rare case that a solution does off deep learning personalization, consider the following pros and cons:

TRAINING TIME

"Training time" refers to the effort of training the machine to make accurate decisions. Deep learning models require lots of training and handholding to get up and running. To get a deep learning model to start delivering accurate decisions requires extensive manual fiddling. Due to the limited number of classifiers, and the black-box nature of deep learning, training often involves a resource and time intensive combination of intuition and trial and error. Getting deep learning to work is hard.

Contrary to this, reinforcement learning is easier and faster to deploy and, due to the visibility you have over decisions, it is easier to troubleshoot if errors occur. The time spent on training traditional and reinforcement models is relatively short.

PROCESSING POWER

The deeper the neural network, the more processing power will be required for training. Yes, this power will be processed by your ML solution of choice, but there's a high likelihood of increased cost being fed into the cost of the contract, and then there is also the environmental impact.

Deep learning networks require significantly more energy but with minimal improvement in effectiveness to typical use cases. As highlighted by Forbes, a recent study shared that *"training a single deep learning model can generate up to 626,155 pounds of C02 emissions – roughly to the total lifetime carbon footprint of five cars"*. And that's just one model.

QUANTITY OF DATA

Deep learning requires a lot of incredibly good data. Because deep learning is a multilayered ML process, it needs a wide variety of inputs to work. The marketing team would need a company-wide effort to create a business strategy to justify a deep learning project.

The saying "garbage in, garbage out" is very applicable here, referring to the need for quality data to inform a model or algorithm, and deep learning is no exception. Data needs to not only be clean and accurate, but also up to date. Only very high-volume situations with remarkably high payouts would justify an investment.

VISIBILITY OVER DECISIONS

Even with excellent data and a high-profile, high-volume problem to solve, deep learning presents risks. Decisions are made in a "black box," so marketers have no transparency as to why certain decisions are made, which makes it hard to build on the outcomes with other programs. Extracting insights and analytics to justify the investment can be hard and reporting to other parts of the organization is also difficult. Essentially, deep learning means that the machine makes decisions, and the marketers are hands-off.

There is also the risk of putting bias in the system, even with good data, which could derail outcomes. Knowing that the right data, or algorithm was implemented is difficult as the machine quickly evolves and makes additional changes.

EASE OF INTERNAL ADOPTION

Even if you have the world's best ML models, with short training times, high visibility reporting, etc., your ML engine will be useless if your team can't make full use of its power. For this reason, ease of adoption is the most important evaluation criterion in this list.

Incremental, approachable implementations of ML allow marketers to gradually improve their processes and provide results and visibility at the same time. This helps with buy-in across teams and can help with measurability. When people see that even a small ML project can pay off, they will be ready to partner on cross-functional and cross-channel projects – for example, combining the ecommerce and marketing experience with shared data and ML-driven promotions.

Ease of adoption comes from visibility over decisions and lighter lifts in terms of implementation, but ease-of-use and flexibility are also key factors. The interface in which your ML engine is operated needs to be intuitive and the platform needs to be flexible enough to adapt to whichever use case you need, whether you are personalizing content, pricing, search outputs, social proof messaging, etc.

The UI and flexibility depend on the personalization solution, but deep learning does not facilitate internal adoption of ML in general, as it does not provide visibility, reporting, or clear “before and after” results. With the length and complexity of a deep learning project, marketers risk losing the interest of other teams, who may decide to go in a different direction.

	Reinforcement Learning	Deep Learning
TRAINING TIME	*Short*	*Long*
PROCESSING POWER	*Small*	*High*
AMOUNT OF DATA	*Small*	*Big*
VISIBILITY OVER DECISIONS	*High*	*Low*
EASE OF INTERNAL ADOPTION	*Medium*	*Difficult*

Questions To Ask Your Next Personalization Solution

To help discover if the solution in your evaluation sights is well suited to you, here are some questions you should ask to ensure their ML is a good fit for your business.

Q. How long will it take for the machine to start making intelligent decisions / how much data is required?

You don’t want to invest in a solution that requires you to get all of your data aligned before you can start seeing returns. Ensure that you can leverage out-of-the-box targets alongside your own data and get a rough timeline for when the machine will start to make intelligent decisions.

DESIRED ANSWER:
“Pretty much no time at all! Even if you have no customer data, you can leverage out-of-the-box behavioral targets, and then feed in your first- or third-party data when the time is right. Our ML will start learning and making intelligent decisions from day one.”

Q. Can I see example case studies of your ML delivering results for other brands / companies?

This question will uncover whether notable brands and marketers are utilizing and trusting the solution's ML. Be sure to check the recency of the customer stories, and whether the solution has the backing of notable brands. You can also ask for customer references.

DESIRED ANSWER:
"Yes, here are our top 5 case studies."

Q. How granular can I go when looking at why the ML has made certain decisions?

Be sure to drill into this point. Will you have visibility over decisions, both for troubleshooting purposes and for developing deeper insights on your customers?

DESIRED ANSWER:
"We're no black-box – you can see which data points have influenced the decisions of the ML engine. You can also access in-depth insights to use across other campaigns and initiatives or to deepen your understanding of your customers."

Q. How easy is it for teams to use? How much of a heavy lift shall I expect in terms of resources and training?

The phrase you're looking for here is 'Marketer-driven.' It doesn't sound all that easy to use if you need a team of developers, data scientists, project managers, merchandisers, marketers, etc., to get a few experiences live.

DESIRED ANSWER:
"'Ease-of-use' and 'ML-Powered' are two of our three product pillars. We have designed the platform to be machine-powered and marketer-driven, meaning that a single marketer can launch and understand the learnings of powerful personalization experiences. Training for this aspect of the platform is typically covered in a single 90-minute training session."

Q. Are there limitations on where I can use ML? What about pages with minimal levels of traffic?

If you don't have the flexibility to use the technology on the channel or the use case of your choosing, you should look elsewhere.

DESIRED ANSWER:
"Our ML sits within a complete platform in which marketers and merchandisers can test, segment or 1-to-1 personalize a range of 'experiences.' An 'experience' can include anything from messaging, content, product recommendations, product badging, social proof messaging, search outputs, and more.

You're not just limited to web and mobile; SDKs (Software Development Kits) and Server-side options makes it easy to deliver ML powered experiences to an array of frameworks and applications or channels, from a React SPA (Single Page Application) through to an in-store kiosk."

Evaluation Criteria

If all the above points are covered, a deep learning model can deliver more accurate decisions in more complex situations as long as you're looking at high-quality and accurate data. The need for heavy development and level of scalability, however, outweighs the improved accuracy.

The truth is deep learning and reinforcement learning are not mutually exclusive. However, starting with a deep learning project can derail overall digital transformation, as it sucks up resources and operates in a black box, with only marginal improvements in outcomes except for the most special situations. Better to start with more accessible ML like reinforcement learning to help get buy-in after you see ROI.

A good ML engine, whether it's deep or reinforcement learning, should help you quickly scale the process of exploitation and exploration, alongside allowing you to understand the results - Balancing learning with performance.