Gartner for Marketers

Gartner

THREE QUESTIONS MARKETING LEADERS SHOULD ASK BEFORE APPROVING AN AI PLAN

(Andrew Frank)

SUMMARY

Artificial intelligence fever is invading the marketing practice. Marketing leaders charged with employing emerging capabilities to grow their business need a practical method of assessing the risks and benefits of AI proposals.

Key Challenges

- As excitement builds around artificial intelligence (AI) applications in marketing, marketing leaders are under pressure to judge whether to sanction potentially risky AI endeavors.
- AI algorithms can exhibit both uncanny predictive accuracy and weirdly inhuman behavior, defying intuitive understanding or explanation and making them hard to trust.
- Real-time decision making on a personal, individual level is a key focus of AI applications. These have the power to greatly enhance customer experience, but also carry legal, ethical and brand risks that are not often visible through a standard cost-benefit lens.

Recommendations

Marketing leaders responsible for marketing technology and emerging trends should:

- Classify automated decisions about offers and experiences by the magnitude of risks posed by advanced profiling and unpredictable customer interactions.
- Distinguish AI algorithms based on transparency: Can they explain their decisions?
- Evaluate the feasibility and costs of monitoring AI decisions with humans in real time to address risks of full automation.

AI applications have the power to greatly enhance customer experience.

Published: 12 July 2017
ID: G00317268

Introduction

The rapid rise in hype and promise surrounding artificial intelligence applications in marketing puts marketing leaders in a bind. Confronted with a growing chorus of internal and external requests for approval of all kinds of AI projects promising marketing breakthroughs, marketers need guidelines to accelerate decision making and manage risk without crushing innovation. Acting too aggressively can unleash unpredictable forces, like Microsoft's former Tay chatbot that went shockingly rogue after a few hours of exposure to abusive dialogue on Twitter in 2016.[1] Acting too conservatively poses the risk of being left behind in a fast-moving revolution that may be transforming the competition and the market's expectations.

The framework in Figure 1 outlines three key questions to help you assess the risk associated with any proposed AI-based automation system.

Figure 1. Three Questions for a Marketing AI Risk Analysis Framework

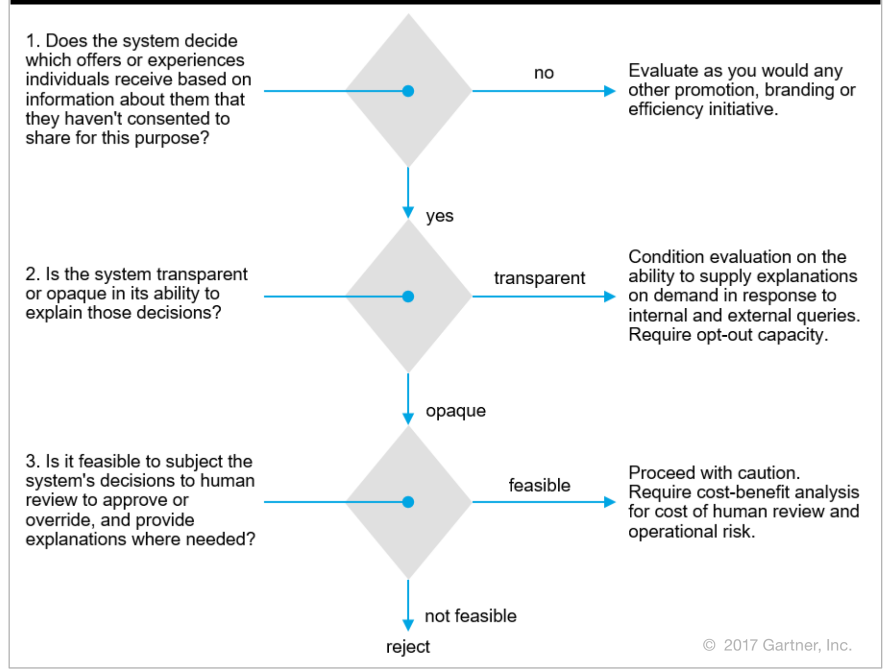

Source: Gartner (July 2017)

ANALYSIS

Classify Automated Decisions About Offers and Experiences by Magnitude of Risk

The AI revolution touches many aspects of marketing, with varying levels of risk. Figure 2 shows the relative positioning of three popular AI marketing applications:

- **Marketer-facing analytics.** Designed to interpret and present complex analysis in simple, comprehensible terms — often in natural language. Examples include Lucy, an IBM Watsonbased system from Equals 3 and Adobe Sensei's natural-language contribution analysis (in development). Although there's always a risk of misinterpretation, such tools are relatively low risk and can easily be contained. Although marketer-facing analytics are valuable, they tend to duplicate some of the explanatory services data scientists provide, making their adoption less transformational than other applications and invisible to outsiders.
- **Untargeted conversational agents.** Designed to offer consumers or business customers access to natural-language dialogs, often through an intermediary such as Amazon Alexa, Google Assistant or Facebook's chatbots. Although in some instances you can tailor these to a known, consenting individual, many are untargeted and scripted using a limited vocabulary and basic decision rules that don't require artificial intelligence. While poor design or execution can lead to a frustrating user experience, legal risks are low as long as profiling is avoided. Still, by providing a new way to interact with your brand, these agents can have a significant impact on customer experience.
- **Real-time personalization.** Designed to automate decisions about what content or offers to present to a customer based on each individual's unique profile, behavioral history and contextual circumstances. The use of individual data to drive algorithmic decisions raises a number of potential risks, as described below. You should thus isolate this class of AI systems for further risk evaluation, as shown in Figure 2.

Figure 2. Common AI Marketing Applications

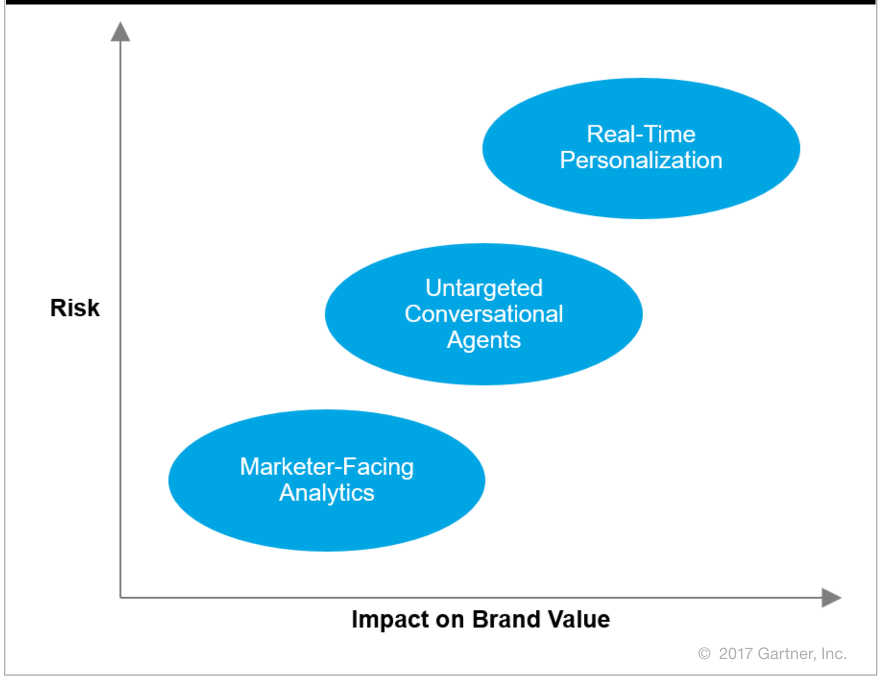

Source: Gartner (July 2017)

Real-time personalization itself presents a hierarchy of risks. At the high end, decisions about which offers or prices to present to which customers based on segmentation can have legal and ethical repercussions, particularly in regulated industries. Even unregulated industries must abide by antidiscrimination laws, which vary by state in the U.S. and by country.

An AI algorithm can predict a surprising number of personal details — including membership in a legally protected class — from a remarkably limited number of clues. Figure 3 illustrates this point with data from a 2013 study that demonstrated how Facebook Likes could predict a range of sensitive personal attributes, including sexual orientation, race, and religious and political views.[2]

An AI algorithm can predict a surprising number of personal details.

Figure 3. Prediction of Personal Traits Based on Facebook Likes

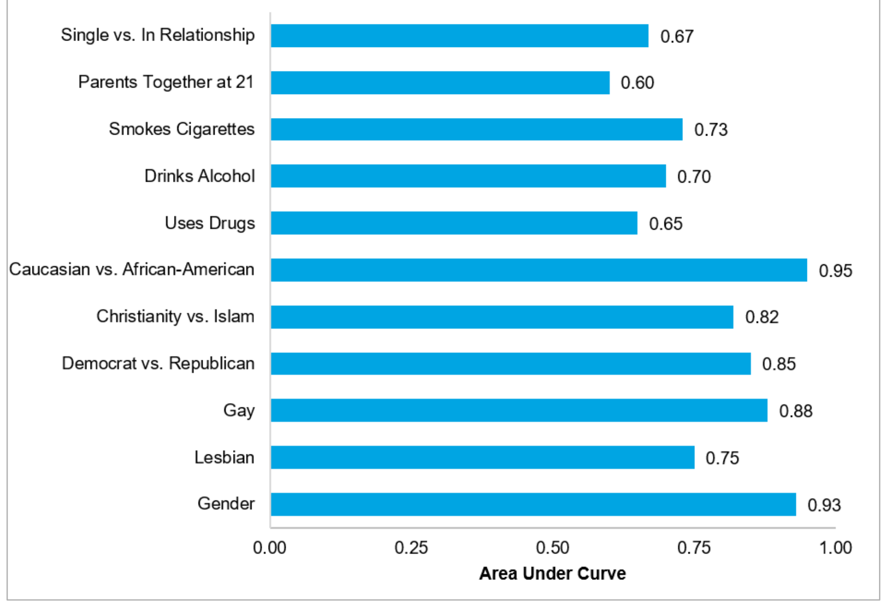

Area Under Curve = cumulative success rate of predictions (statistical accuracy)

Source: PNAS Online, National Academy of Sciences[2]

This inferential power has several worrisome implications. In Europe, Section 4 of the General Data Protection Regulation (GDPR), which goes into effect 28 May 2018, specifically grants individuals the right to object to and obtain exemption from "automated individual decision making."[3] Even beyond the jurisdiction of GDPR, however, marketers must consider the implications of profiling consumers who may object to the practice in principle, without explicit consent. Even where data is collected anonymously and in aggregate, profiling for personalization can go awry.

For example, in 2016 Amazon had to redress a scandal over the revelation that its same-day delivery service was unavailable to certain ZIP Codes in predominantly black neighborhoods.[4] Such stories illustrate how algorithmic segmentation can result in unintentional systemic inequality or inadvertently expose a protected class.

At the lower end of the risk hierarchy are situations where techniques such as automatic content generation are used to optimize language or imagery in targeted advertising. Systems such as Persado and Wylei can generate, test and tune messaging to optimize audience response, but marketers can be surprised when messages appear off-brand or even incomprehensible to them. Providers generally include tools for marketers to control which portions of a message or assemblies of a canvas are immutable, but the perceived riskiness of such techniques is largely dependent on corporate culture.

An AI project that presents no customer-facing risks may have other risk factors you need to consider, but you can generally treat these as you would any other proposed marketing initiative. However, be sensitive to second-order effects: for example, algorithms affecting decisions that appear removed from customer experience, such as distribution optimization, may have unintended customer impacts.

Distinguish AI Algorithms Based on Their Transparency

Identifying an artificial intelligence proposal as potentially risky is not necessarily grounds to reject. Instead, it indicates a need to establish confidence that the system will behave within the bounds of acceptable risk.

Among the many ways to classify AI algorithms, one helpful distinction in this context is transparency.[5] Transparent algorithms can explain their analysis and decisions — in the most advanced cases directly in natural language, but in others with the help of a human analyst. Various forms of regression fall into this category. You can think of these algorithms as being on the "shallow" end of AI, in contrast to "deep" methods that are attracting considerable attention today.

At the other extreme, opaque algorithms like deep learning can achieve major conceptual breakthroughs, such as beating the world champion at the ancient Chinese board game Go. But their decision process is buried in the so-called "hidden layers" of a deep neural network (DNN). Explaining the decisions of a DNN is a subject of advanced research, but despite their abilities the behavior of DNNs can be counterintuitive and unpredictable.[6,7,8,9] Some DNNs, such as those use in image recognition, can display intermediate results (such as nodes that recognize edges or faces or facial features in an image). But when the subject of analysis is more abstract, such as categorizing features of human behavior, inferences based on deep learning may be less intelligible.

Nevertheless, making explanations of profiling available on demand to staff and customers should be a condition of deployment whenever feasible. Given the difficulty inherent in explaining algorithms, many software providers put the onus on clients to, at minimum, reveal the data that contributed to the decision without providing details about the process. This is both mandated by GDPR and an ethically sound practice in general, along with offering consumers opt-out, deletion and correction choices relating to personal data wherever they're exposed to profile-based personalization.

Deep learning, like many other AI algorithms, depends on data to learn, so confidence in the quality of your training and testing datasets is essential before getting started. However, before you put an opaque system into production, you'll need a thorough quality assurance (QA) process to review and validate its decisions. If the risks warrant it, you may need an ongoing review process in your production environment, at least until you can establish a dependable track record. But be aware that for certain algorithms, as learning continues results may deviate from established expectations and the smartest algorithms never stop learning.

Evaluate the Feasibility and Costs of Monitoring AI Decisions With Humans

If you're considering an opaque algorithm that makes substantive decisions about customer experience, your final question must be whether you can cost-effectively validate it with human review. Ideally, every decision should be subject to human review, although this may not always be feasible at scale. Figure 4 illustrates a model QA process for learning systems that require training, such as a DNN.

The training, testing and monitoring process supports two levels of feedback once a system is trained. At the tester level, testers can both simulate personalized user experience and flag "exceptions" — that is, experiences that don't align with expectations. Testers also can intervene in production system decisions, replacing a problematic experience with a generic alternative. At the audience level, "objections" refers to signals users may generate, such as clicking on a "hide ad" or "privacy policy" link to opt out, that suggest the prediction that the customer will appreciate the experience has failed.

Figure 4. AI System Training, Testing and Monitoring Process

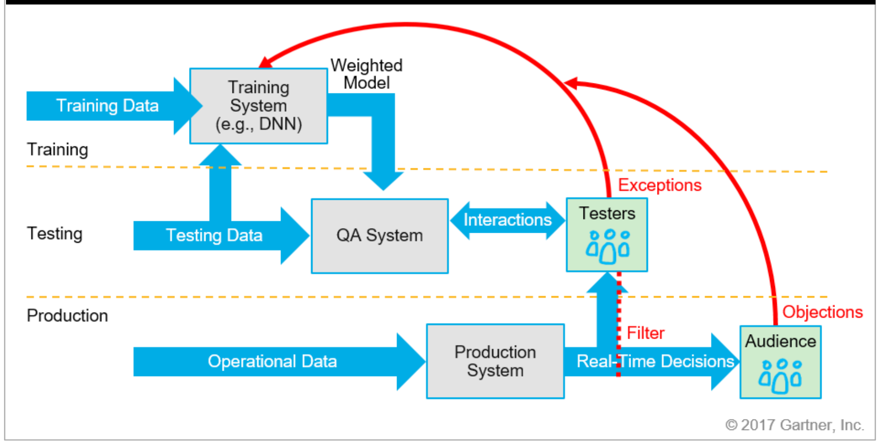

Source: Gartner (July 2017)

Chances are that adding extensive testing and monitoring to an AI proposal will substantially increase its cost and time-to-value relative to the initial proposal, and may render a project uneconomical. The requirement should illustrate why AI projects need to produce transformative economic results — ironically to justify the human resource costs of building trust in their behavior. If the anticipated results can't justify this expense, you should probably reject it. Over time, industry solutions will likely expand to decrease this cost.

GARTNER RECOMMENDED READING

"Cool Vendors in Artificial Intelligence for Marketing"

"Algorithmic Marketing Essentials"

"Bandits at the Gate: Automated Alternatives to A/B and Multivariate Testing"

"Use Natural-Language Technologies to Converse With Audiences at Scale"

"Crawl, Walk, Run: Define Your Vision and Roadmap for Personalization"

"How to Apply Artificial Intelligence to Digital Commerce"

Some documents may not be available as part of your Gartner subscription.

Evidence

This research was developed based on primary and secondary research including interviews with brands and experts as well as inquiries with Gartner clients throughout 2016 and 2017.

1 "Microsoft Is Deleting Its AI Chatbot's Incredibly Racist Tweets." Business Insider. More recently, another Microsoft chatbot, Zo, exhibited similar controversial behavior: "Microsoft Chatbot Turns Controversial Yet Again: Zo AI Calls the Qu'ran 'Violent' and Gives Its Opinion on Osama Bin Laden in Bizarre Chat." Daily Mail.

2 "Private Traits and Attributes Are Predictable From Digital Records of Human Behavior." PNAS Online. National Academy of Sciences.

3 "General Data Protection Regulation." Council of the European Union. Page 146, Paragraph 2: "Where personal data are processed for direct marketing purposes, the data subject shall have the right to object at any time to processing of personal data concerning him or her for such marketing, which includes profiling to the extent that it is related to such direct marketing."

4 "Amazon Doesn't Consider the Race of Its Customers. Should It?" Bloomberg.

5 "AI in Customer Engagement: Balancing Risk and Reward." Pegasystems. Rob Walker, Ph.D., vice president of decision management and analytics highlights the distinction between transparent and opaque algorithms in this video presentation.

6 "Making Computers Explain Themselves." MIT News. Massachusetts Institute of Technology.

7 "Automated Cybersecurity: Removing Humans From the Equation." Nextgov. (Download required.)

8 "Attentive Explanations: Justifying Decisions and Pointing to the Evidence." Cornell University Library. (Download required.)

9 "Deep Neural Networks Are Easily Fooled: High Confidence Predictions for Unrecognizable Images." Cornell University Library. (Download required.)