OfferFit

AI decisioning for the enterprise

Next best ~~action~~ everything

Using AI decisioning for 1:1 personalization

George Khachatryan, Ph.D.
Co-Founder & CEO

Victor Kostyuk, Ph.D.
Co-Founder & CTO

Nathaniel Rounds, Ph.D.
Head of Product Marketing

OfferFit
June 2024

Contents

Executive summary

Marketing has always been awash in buzzwords, and two that have been buzzing the loudest lately are **personalization** and **AI**. Everyone and everything, it seems, is helping marketers personalize, no doubt with the help of AI. In this maelstrom of hype, marketers understandably are looking for simple explanations and clear value propositions. Of course, "AI hype" in marketing is not new. And in fact, most enterprise marketers have been using **machine learning (ML)** to personalize for years.

One common method of using AI models to personalize are so-called **next best action (NBA)** models. On its face, it's a simple idea — a next best action model should tell you the best action to take next with each customer. Who could say no to that? In practice, "next best action" is not a precise term, nor one used by data scientists or experts in machine learning. And unfortunately, these models are often a "black box" for the marketers that are using them. In this paper, we will try to open this black box to explain what NBA models are really doing, their limitations, and the newly emerging state of the art: **next best everything**.

Though they may be a black box, many businesses use traditional "next best action" models with great success, and perhaps 5 years ago, we could say these models were "state of the art." But marketers are realizing that traditional NBA models have significant limitations.

1. **Traditional NBA models are brittle** — keeping them up to date as markets and customer behaviors change is both slow and expensive.
2. **Traditional NBA models typically make recommendations for segments, not people**. They offer personalization which is not very personal.
3. The only way to know if NBA models are really making optimal decisions is through **manual testing** — usually A/B testing or multivariate testing by segment.
4. The "action" the model recommends is usually a product or offer. **Traditional NBA models cannot simultaneously optimize all the decisions marketers must make** with every customer touchpoint — e.g. message, creative, financial incentive, channel, time, day, frequency — much less handle the interdependencies between these decisions.

Savvy marketers are now turning to a newer technology, **AI decisioning**, based on a type of machine learning called **reinforcement learning (RL)**. AI decisioning offers fundamental advantages over the older "next best action" approach.

1. AI decisioning agents **continuously learn and adapt**. They are flexible rather than brittle, and automatically adjust to market changes.
2. AI decisioning agents leverage **all first-party data** about every customer characteristic to make **1:1 decisions** about individuals, not segments.
3. AI decisioning agents **empirically discover** the best decisions for each customer which maximize the marketer's KPIs — **no need for manual testing**.
4. AI decisioning agents can optimize across **multiple dimensions simultaneously**.

AI decisioning agents learn the best decisions to make for each individual customer — not only the right product to offer, but the right time to send the offer, the right day of the week, the right message, the right channel, and the right frequency of communication. AI decisioning is not "next best action" — it is truly **next best everything**.

Traditional next best action models

"Next best action" is not a precise or technical term. In principle, any **business rule** which takes **customer data as input**, and recommends an **action as output**, could be called "next best action." In practice, the term is used somewhat more narrowly. Typically, an NBA model combines **predictive models** and **business rules** to predict the product or offer to suggest to a customer. Let's unpack how that works. A traditional NBA model — which from now on we'll just call an NBA model, though the term could be used in other ways — consists of the following steps.

1. Start with raw **customer data** — e.g. purchases, web browsing events, app logins, email clicks.
2. Use these data to calculate **customer features** — e.g. purchases in the last 90 days, average value of last 10 purchases, days since last purchase.
3. Use these features as the input for **predictive models** — e.g. churn propensity, propensity to repurchase in a given category, customer lifetime value prediction.
4. (Optional) **Segment** customers based on these features and models — e.g. create a "high risk" segment of customers with a churn score over 80.
5. Write **business rules** for which product to each customer based on the features and models - either for the customer population as a whole, or separate rules for each customer segment.

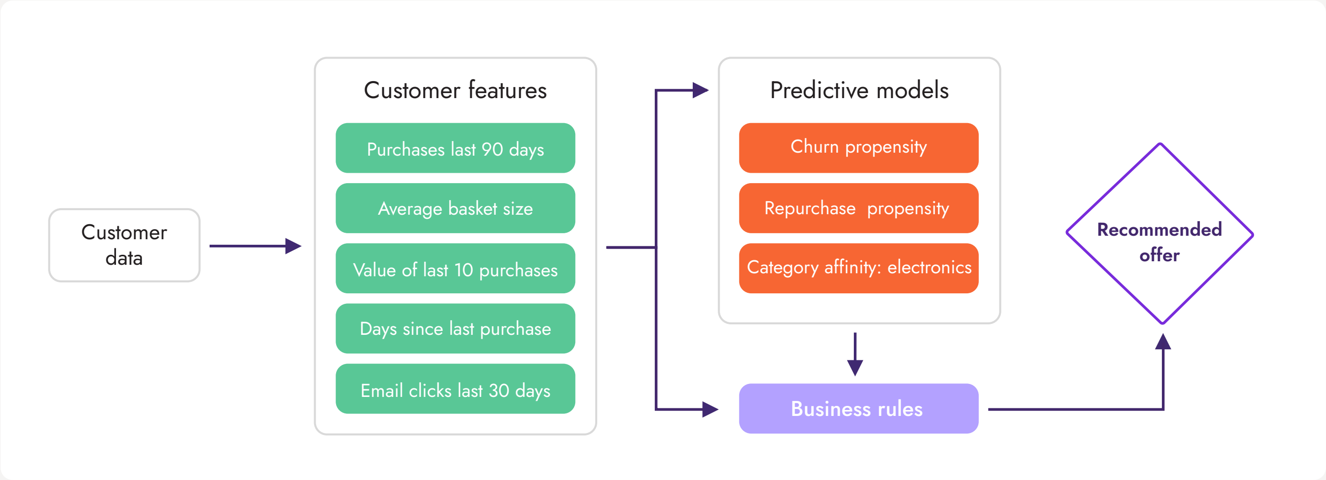

A next best action model

An NBA model combines machine learning — the **predictive models**, based on a type of ML called **supervised learning** — with manual **business rules**. The predictive models could be, for example, churn propensity or category affinity. A churn model might give every customer a score on a scale from 0 to 100 based on how likely the model thinks the customer is to churn in the next 3 months. A category affinity model might predict how likely a customer's next purchase is to be in a given category.

The outputs of the predictive models are then used to define a collection of "if-then" **business rules**. For example, a rule could be: if a customer has a churn propensity greater than 75, then send them a coupon for 20% off to renew their subscription. An NBA model — consists of this collection of customer features, predictive models, and business rules which are used to make recommendations.

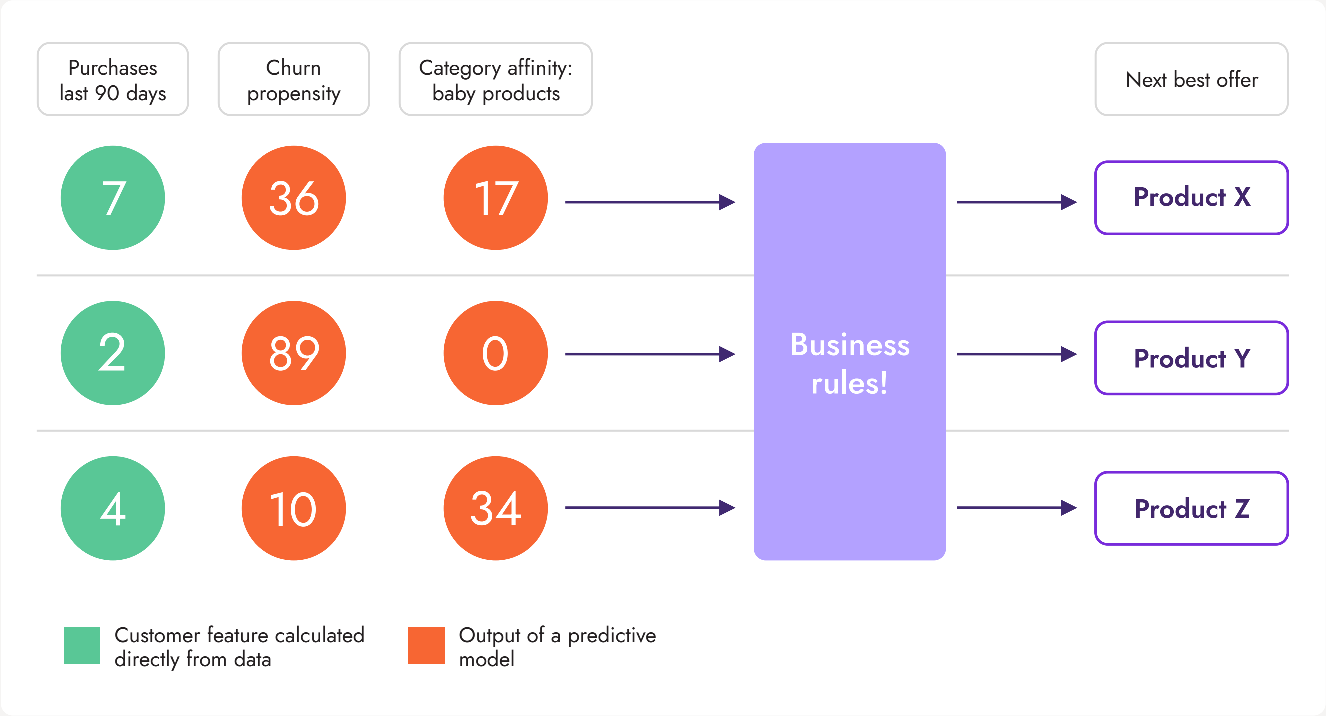

A "next best action" model uses business rules to recommend product offers. Each row represents a customer. Note that the first input is a simple data point, while the others are outputs of predictive models.

How do we know an action is really best?

The models we've described can certainly be called "next action" models — they tell a marketer what action to take next with each customer. But why do they deserve to be called next *best* action? Suppose an NBA model says that a particular pair of jeans is the "best" next product for a particular customer. What does that really mean? Here's the most likely sequence of events:

1. The NBA model includes category affinity models for various product categories. The customer's highest affinity score is for the "jeans" category.
2. The NBA model's business rules say "recommend a product from the category for which the customer has the highest affinity score."

In this case, the NBA model is predicting that these jeans are the product the customer is most likely to buy next. But what the marketer needs to know is what action to take to maximize the metrics that matter to the marketer — in this case, perhaps CLV, repurchases, or revenue.

Suppose a bank has an NBA model to recommend products to offer their customers. This ML model might predict that a given customer is most likely to apply for a particular credit card. Is it a good idea to send that customer offers for the card? Unfortunately, the NBA model doesn't answer that question. Any of these situations is possible:

- The customer would apply for that credit card without any prompting, so it would have been more valuable to email the customer about something else.
- The customer is most likely to apply for that card if the marketer takes no action, but would also be responsive to offers for some higher margin product such as a premium credit card.
- Marketing communications will increase customer engagement, but won't end up influencing the purchasing decision at all.

Knowing what customers are most likely to do **without an intervention** doesn't tell marketers what **intervention would be most effective**.

This gap between the model's prediction and the optimal marketing action can be invisible. Perhaps a model predicts that a customer is likely to buy a particular product, so the marketer sends the customer an incentive. Lo and behold, customers who got the incentive bought the product more often — the system works! The model has introduced **confirmation bias**. No one says no to a discount for a product they were already planning to buy, but those same customers might have bought a higher-margin product, or accepted a less generous incentive, if the marketer had sent those offers instead.

A "next best offer" prediction that a customer is likely to buy a given product doesn't necessarily mean that offering that product, or indeed taking any particular action towards that customer at all, is the action most likely to maximize the metrics which matter most to the business.

Segments and manual testing

For a "next best action" model to be of any real value, marketers need to know that the actions the model is recommending really are the ones with the most significant impact on their business KPIs. The typical solution goes like this:

1. Make the business rules as "personal" as possible: divide customers into microsegments and have different rules for each segment.
2. Run A/B tests or multivariate tests to validate the business rules for each segment.

Marketers can segment customers based on the outputs of their predictive models, perhaps into quintiles or deciles. Combining the output of various models will give perhaps hundreds of **microsegments**. For example, suppose a business picks 4 customer characteristics as the key dimensions by which they wish to segment their customers. The business might split customers into five segments for each — customers with churn scores between 81-100, 61-80, and so forth. Four dimensions, each split into gives segments, gives 5*5*5*5 = 625 microsegments.

Dividing customers into quintiles using two predictive models — churn propensity and affinity for baby products — gives 25 customer segments.

Next, marketers need to validate that their NBA model is optimized for each segment. Running simultaneous A/B tests on 625 segments is likely infeasible, so businesses typically use **multivariate testing** to measure what works best with each segment. From there, marketers update their **business rules** about which actions to take for each segment.

The result, hopefully, is a "next best action" model which gives a personalized prediction of the best offer to make every customer. This method can be effective, and indeed is how many enterprises use ML today to personalize. Perhaps 5 years ago, this method was "state of the art." But this approach is not without its limitations.

Next best action models are brittle and expensive to maintain

A next best action model is built with predictive models that are trained and deployed at a point in time, but markets change, and customer behaviors change with them. Suppose a business has a well-built churn prediction model trained on existing data, and business rules based on extensive multivariate tests. Next week, a marketer has an idea for a new offer, creative, or incentive to send to customers at risk for churn. These offers are new — the marketer's multivariate tests did not include them, and so their business rules have no data about them. And how will these new offers impact customer churn propensity? Both the marketer's business rules, and the underlying predictive models and the customer segments they define, may begin to go awry.

The business now faces a choice: either data scientists must collect training data on the new offers and retrain their models, or marketers must run the risks of bad predictions from an out-of-date model. Simply put, predictive models are brittle. They work well when applied to data very similar to the data they were trained on, but have trouble generalizing to new situations. It does not take a seismic shift like a pandemic or a recession to make the model run off course. Deploying a new offer, or emailing customers with a different frequency, is enough for the models' predictions to become unreliable. As the collection of models a marketer deploys — churn prediction, repurchase prediction — continues to grow, the effort and expense to keep those models up-to-date will grow with them.

Even if businesses devote time and money to continuously retraining their models, marketers aren't out of the woods. After all, the marketer ran multivariate tests at a moment in time. The business rules based on the results of these tests are now out of date. The marketer can retest, but that's more time and expense. Meanwhile, markets and customer behaviors continue to evolve. The business's choices are not very appealing: devote significant resources to continuously retrain models and rerun tests, stick with outdated segments and rules and hope for the best.

Personalization by segment is just not very personal

Even if businesses spare no expense in updating their models and retesting their business rules, there is no escaping a more fundamental problem. "Segments and rules" based personalization is just not very personal! 625 microsegments may sound like a lot, but when a business has thousands or millions of customers, microsegments are a long way from 1:1 personalization.

Marketers want to find the "winning offer" for each microsegment. The results of multivariate tests for 625 segments certainly seem very personal. "For segment X, offers for product Y converted more than offers for product Z." Sure sounds like a win! The problem is that each microsegment represents hundreds or thousands of customers, and some of those customers preferred product Z. If marketers adopt the "winning" option, they are simply letting the majority rule. In an election, getting 51% of the vote means you win. In a marketing campaign, engaging 51% of your customers is probably a losing strategy — 49% of your customers aren't reading your email. Dividing the customer base into 625 microsegments merely repeats the problem 625 times.

Let's look at a concrete example. RetailCo uses a next best action model to send product offers to their customers. We'll consider two customers in one microsegment, and the model's recommendations for each of them.

The orange segment represents customers with baby product affinity between 61-80, and churn score greater than 80. (RetailCo uses four customer dimensions to define microsegments; for simplicity the diagram shows only two of them.)

Benjamin and Michael are two of the thousands of customers who belong to the orange microsegment in the diagram above. They used to buy a lot of diapers from RetailCo, but they haven't repurchased lately. RetailCo's next best action model recommends we email them a 25% off coupon for diapers. Is that a good idea? To answer that question, we'd need to know more about who these customers are, and why they've stopped buying.

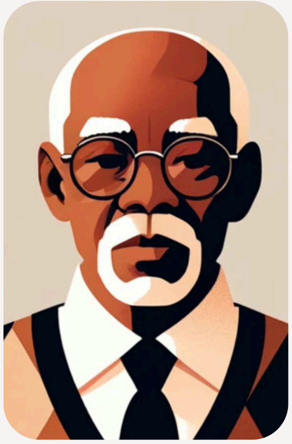

	Benjamin	Michael
Propensity to churn	83	87
Propensity to make a second purchase	23	12
Category affinity, baby products	72	63
Purchases last 90 days	0	0

Benjamin and Michael both used to buy a lot of diapers from RetailCo, but it doesn't take a churn prediction model to see that they haven't been buying lately.

- Benjamin is a grandfather who provides a lot of childcare for his grandchildren. However, he is on a fixed income and is very price sensitive. Though the churn model doesn't know it yet, he's already churned — he found a lower price. The 25% off coupon will inspire Benjamin to buy more diapers from RetailCo. Unfortunately, 25% off means RetailCo is selling at a loss, and it doesn't result in winback. Next month, Benjamin goes back to the cheaper competitor. A better move would be to send him a "subscribe and save" offer — the idea of a recurring discount will be compelling to him, and the subscription means he'll spend less time comparing prices.
- Michael just graduated from college. During his senior year, he babysat for his landlady and sometimes bought diapers for her as a favor. He has since moved to another city, and he won't bother to open an email about baby products. He likes RetailCo; he just doesn't need diapers. What he needs are some shirts to wear to job interviews — he'd happily buy those without a discount.

RetailCo devotes tremendous time and expense to fine tuning and testing their NBA model, but the model's recommendations are missing the mark. Of course, they have much more data about Benjamin and Michael that could have helped them make better decisions. The NBA model simply didn't use it.

Years as customer

Responses to satisfaction surveys

Basket sizes & channels

Purchase frequency

Email interactions

Web browsing behavior

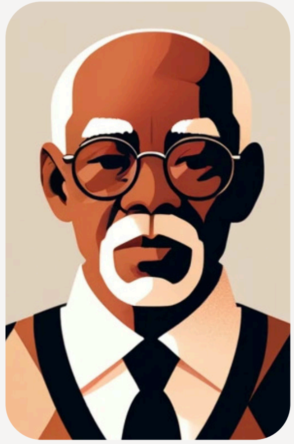

Benjamin

Michael

First-party data that RetailCo's NBA model does not take into consideration.

Marketers' rich first-party data describes hundreds of customer characteristics. Rules built on the results of a handful of predictive models reduces customers to only a few data points — the outputs of each model — letting the depth of first-party data go to waste. Despite the effort businesses are making to continuously update models and rerun tests, the result is personalization that simply isn't very personal.

The need for “next best everything”

We’ve all heard the phrase “next best action” so often that the individual words can start to lose their meaning. Yet as we’ve explained, there are many reasons why an NBA models recommendations might not really be “best.” That begs the question: when is next? And for that matter, what’s an action? These are not pedantic questions — marketers have to make many decisions simultaneously when they engage with their customers.

When marketers speak of an “action,” they most often mean a product, promotion, or offer that they might make to a customer. An NBA model might say that the next best promotion to offer customer X is product Y. But that doesn’t tell the marketer all she needs to know. For example, what’s the right financial incentive? Perhaps the product has been discounted too much, and the customer would have happily converted with less of an incentive.

And that’s not all — the NBA model doesn’t tell the marketer when to make the offer. In other words, when exactly is the “next” in “next best action? Should the marketer email today, or next week? In the morning, or in the afternoon? After all it hardly matters if we offer the right product, if the customer doesn’t read the email. Perhaps the customer is more likely to engage if the email had a different subject line or creative. For that matter, is it best to contact this customer by email in the first place as opposed to a text or a push notification? And if the customer ignores us, how long is it best to wait before contacting that particular customer again?

Clearly, simply knowing the next best product or offer is not enough. While we typically wouldn’t think of channel — or time of day, or day of the week, or the tone of a subject line — as actions, these are decisions that a marketer needs to make. It doesn’t do much good to say “We’re personalizing offers,” if the personalized offers are sent at the wrong time, at the wrong frequency, or over the wrong channel, or with the wrong incentive. And these dimensions are interdependent! Customers read emails at different times on different days. A customer might be interested in buying a new computer monitor at work at 9am on Tuesday, but jump on a discount for a new phone as they scroll at 7pm on Saturday. Sending the same offers at the opposite times would be ineffective.

In principle, a business could have a separate “next best action” model for each decision a marketer needs to make — product, incentive, channel, time of day, day of the week, frequency. But this approach is not viable in practice. The customer characteristics most critical for each decision are different, so the microsegments for each dimension would need to be different. Setting up business rules for all these microsegments, let alone validating them with multivariate tests, is not feasible. Even if it were possible, it wouldn’t solve the problem of interdependence between decisions along different dimensions.

The “next best action” approach is fundamentally limited. Marketers don’t simply need to find the next best **action** — they need to find the next best **everything**.

Why AI decisioning is next best everything

As we have seen, the older "next best action" methods have fundamental limitations. A new approach called **AI decisioning**, based on a type of machine learning called **reinforcement learning (RL)**, offers marketers a true "next best everything." AI decisioning can find not only the next best product to offer every customer, but the best channel, time of day, day of the week, frequency, message, creative, or indeed any other dimension a marketer wishes to test.

An AI decisioning agent **learns** from its environment — it iteratively chooses from a set of actions, receives some reward in response, and then chooses another action based on what it's learned. As the model learns from the rewards it receives, it updates its policy for selecting future actions with the goal of maximizing the reward achieved over time.

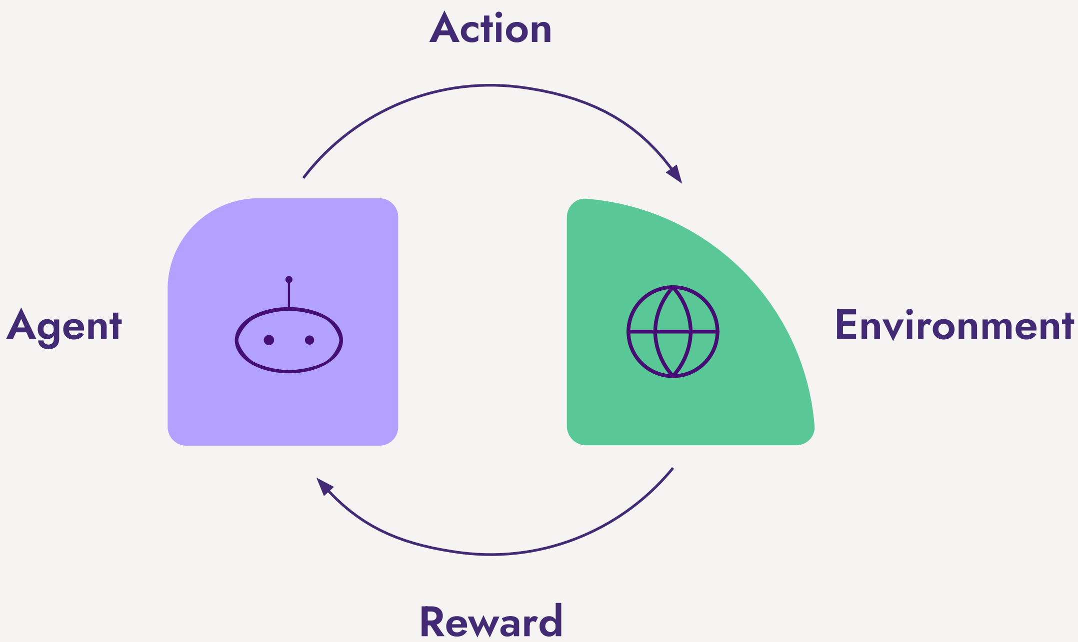

The advantages of AI decisioning agents built on reinforcement learning are fundamental.

- AI decisioning maximizes the metrics that matter to the marketer. There is **no need for manual A/B tests or multivariate tests** to see if the agent is working — the agent itself measures (and improves!) its own effectiveness.
- AI decisioning agents **continuously learn and adapt**. They are flexible rather than brittle markets or customer behavior changes — the agent learns and adapts autonomously.

- As a result, AI decisioning is **much easier to set up and maintain** than traditional NBA models — a single agent, which continuously learns and adapts, can replace an intricate system of models, rules, and manual testing.
- AI decisioning agents can use **all first-party data** about every customer characteristic to make **1:1 decisions** about individuals, not segments. They are not constrained to only consider a few data points or the outputs of a few models.
- AI decisioning is **additive to a business's existing predictive models**. The outputs of these models can be inputs to the AI decisioning agent. (Predictive models aren't necessary for AI decisioning to work, and it's better to start with AI decisioning. But if a business has predictive models already, they will make AI decisioning even better.)
- AI decisioning agents **empirically discover** the best 1:1 decision for each customer — they experiment, rather than relying only on what's worked in the past.
- AI decisioning optimizes across **multiple dimensions simultaneously**, finding the right choice of product, message, timing, day, etc., even when these choices depend on each other.

Traditional next best action models are valuable, and still get good results. But we can no longer say they are state of the art. AI decisioning offers marketers what they really need — not next action, but next best everything.

About OfferFit

OfferFit is an AI Decisioning Engine for lifecycle marketers.

Our AI decisioning agents make 1:1 decisions on the optimal way to market to each individual customer. These agents start with a goal — a business KPI like incremental revenue or CLV — and a set of actions they are allowed to take – options for channel, message, product, creative, time, day, and frequency of communication. The agents then use reinforcement learning (a type of machine learning) to make decisions for each customer, aiming to maximize the goal. AI decisioning agents autonomously experiment and continuously learn from every customer interaction. OfferFit works with top brands in telecom, energy, retail, travel, streaming video, and financial services, among others.

OfferFit

AI decisioning for the enterprise

About the Authors

George Khachatryan
Co-Founder & CEO

Before founding **OfferFit**, George was an Associate Partner at McKinsey & Company, where he led transformations of some of the world's largest companies. Previously, George co-founded a technology company whose products are used today by over 1 million users. He holds a Ph.D. in Mathematics from Cornell University.

Victor Kostyuk
Co-Founder & CTO

Before founding **OfferFit**, Victor was a Lead Data Scientist at the Boston Consulting Group. He is an expert in reinforcement learning, with extensive experience serving Fortune 500 companies on offer personalization. He holds a Ph.D. in Mathematics from Cornell University.

Nathaniel Rounds
Head of Product Marketing

Before joining **OfferFit**, Nathaniel spent 10 years leading design and developing of SaaS products, with a focus on educational content. He holds a Ph.D. in mathematics from Stony Brook University.