Baynote

The Human Need for Personalization:
Psychology, Technology and Science

Scott Brave, Ph. D

July 18, 2012

An executive white paper defining the psychology, technology and science that underlie a shopper's desire for a personalized ecommerce experience. This paper looks at the definition of personalization, the psychology of needs, the technological constructs that make automated personalization systems possible and the science behind the man-machine interface that brings user psychology and technology together. This paper is for ecommerce executives interested in improving the customer experience on their site(s) and who want a better understanding of what personalization is and how it plays into the overall customer experience.

Scott Brave, Ph. D

Scott is a founder and Chief Technology Officer of Baynote, Inc. Prior to Baynote, he was a postdoctoral scholar at Stanford University and served as lab manager for the CHIMe (Communication between Humans and Interactive Media) Lab. Scott is an inventor of six patents and co-author of over 25 publications in the areas of human-computer interaction and artificial intelligence. Scott was also an Editor of the "International Journal of Human-Computer Studies" (Amsterdam: Elsevier) and co-author of "Wired for speech: How voice activates and advances the human-computer relationship" (Cambridge, MA: MIT Press). Scott received his Ph.D. in Human-Computer Interaction and a B.S. in Computer Systems Engineering from Stanford University, and his Master's from the MIT Media Lab.

Table of Contents

Baynote, Inc. | 333 W San Carlos St, Ste 700, San Jose CA 95110 USA | | sales@baynote.com | partners@baynote.com | baynote.com

Introduction

We talk a lot about personalization in ecommerce. We want our sites to provide it, and we want our customers to feel it. But what is it exactly? How do we achieve it? What impact does it really have? Personalization, after all, is a combination of many factors — part psychology, part science and part technology — and behind it all lays the magic of human judgment. Like three legs of a stool, it is impossible to talk about personalization without first digging into the psychology of the shopper — their motivations and response mechanisms. At the same time, science coupled with technology allows us to look at the data inherent in personalizing online shopping experiences and to deliver it through a variety of technologies in an ordered fashion. This plays well with both ecommerce executives and shoppers themselves.

> *"Human instinct, intelligence and decision making, are at the core of both the scientific algorithms that drive personalization schema as well as the technology architecture that delivers that schema to the ecommerce system and the shopping experience."*

But as with most things, to err is uniquely human. Human instinct, intelligence and decision making, are at the core of both the scientific algorithms that drive personalization schema as well as the technology architecture that delivers that schema to the ecommerce system and the shopping experience. These factors drive shoppers to behave in predictable and unpredictable ways. The net result is that the myth of one to one personalization is just that — a mythical end state that can never truly be reached, only hoped for. So what does personalization really mean? In this whitepaper we will dig into these big topics and provide you with an understanding of these core personalization frameworks. After reading this paper, I hope that you can truly understand what personalization means and are able to lead rigorous, productive discussions on personalization and related strategies for your company. This paper will provide you with the knowledge you need to ask the right questions.

The Psychology of Needs

To begin, let's get grounded in a definition of personalization. Personalization is a concept that, while sounding simple, has several layers of complexity. We start with the idea that personalization is about giving each individual user a unique experience. What if Fred Smith comes to your ecommerce site and is presented with the perfect product to meet his needs, but a thousand other people are shown the same product and have an identical experience of the perfect fit? Is this no longer personalized? Of course it is. The experience provided Fred with exactly what he needed at that moment. While uniqueness is often a byproduct of personalization, it is not a requirement.

The goal of personalization is:

To provide customers with interactive experiences tailored to optimally satisfy their individual needs.

The concept of "needs" is a critically important one in psychology. Needs — which should be understood more generally as needs, goals, wants and preferences — are the motivating factor behind everything we do. They drive all of our behaviors, including the relationships we cultivate and the objects we consume. Every emotion we have is a result of the degree to which our needs are satisfied or inhibited. Personalization is much like a matchmaking exercise with the ultimate goal of pinpointing the best product or service from your catalog that best satisfies your customers' needs.

The psychological literature provides a number of alternate categorizations of human needs. In general, it is agreed that there is a set of fundamental or basic needs shared by all such as the need for survival, safety, belonging, esteem and self-actualization. These make up the core of Abraham Maslow's hierarchy of needs. His psychological theory was first proposed in his 1943 paper "A Theory of Human Motivation."[1]

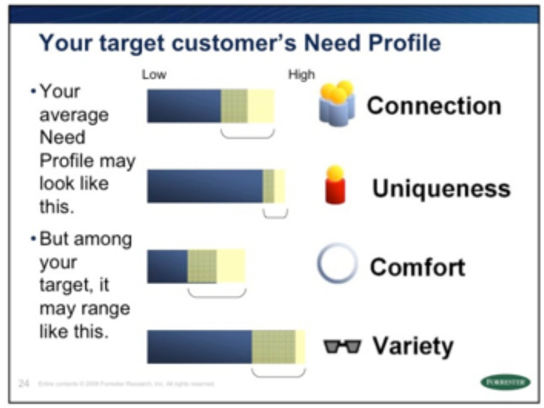

McQuivey's Needs Assessment[4], Figure 1

In 2008, Forrester Research held a conference during which James McQuivey, Ph.D. the foremost analyst tracking and defining the power and impact of digital disruption on traditional businesses, delivered a keynote presentation. In this presentation, McQuivey argued that a more valuable categorization for ecommerce is the universal and at times conflicting need for connection, uniqueness, comfort and variety. According to Mr. McQuivey, everyone has all four, but they vary in importance by individual and can shift over time due to changing circumstances. People will ultimately trade off one need against another[2]. Regardless of categorization, what is perhaps more prescriptive when considering personalization strategy is how these very basic needs crystallize into purchase decisions. Customers generally don't come to your website thinking fundamentally "I'm looking to satisfy my need for survival and uniqueness." More likely, they are thinking conceptually, "I need a cool, retro T-shirt to wear to the outdoor festival," or even objectively, "I want that green and blue Grateful Dead T-shirt I saw in the store." These are all the same need, at different levels of specificity.[3]

The strategy for meeting customer needs looks very different depending on the specificity of the customer's intent. If a customer knows or thinks he knows precisely what he wants, showing him the most relevant product or related set of products is usually best. The worst thing the site can do is throw out random-seeming suggestions that distract from the task at hand. On the other hand, if the customer is just looking for something cool, then effectively meeting that need may be better approached from the perspective of discovery. Discovery is more about delight and surprise in stumbling upon products which meet a customer's more abstract conceptual or fundamental needs. Relevancy tends to be more utilitarian in nature and focused on more concrete conceptual and objective needs.

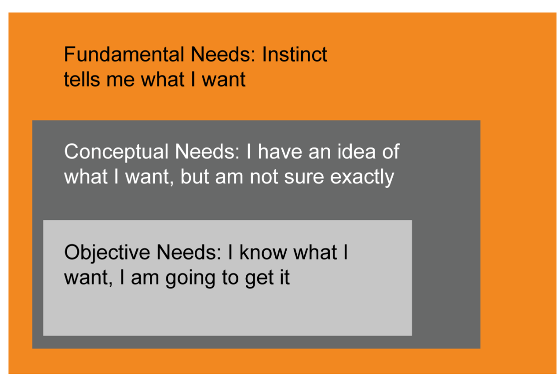

The Need Nest
Figure 2

The trick is pinpointing your customers' intent somewhere along that relevancy-discovery continuum. Certain ecommerce sites may lean more towards one end of the continuum or the other as a general rule. For example, compare shopping for tools and electronics to apparel and media. As a shopper, my mode may also change as I move through the buy flow on a particular site. Someone initially looking for a cool gift may discover the idea of buying a coffee-related gift, and gradually narrow his scope to specifically look for a compact, mid-priced espresso machine. In this framework, taste can be seen as a higher-level conceptual/aesthetic need from which multiple objective needs emerge over time.

The Consumer Buy Flow

Figure 3

While we are on the topic of needs, I would be remiss if I did not mention the other half of the equation: your needs as a retailer. The best product for your customer may not be the best fit for your business. For example, what if the second best product match for your customer's need is actually the best product for you from a profitability standpoint? Which product should you suggest first? Should you promote higher-margin products? These questions highlight the tension between creating personalized customer experiences and creating business value to meet the needs of both parties. This second view of needs helps to distinguish between the techniques of personalization and targeting.

True personalization approaches matchmaking primarily from the side of the customer, aiming to find the best fit for his or her individual needs. Conversely, targeting approaches matchmaking primarily from the business side, looking to reach customers who are most likely to accept a specific product or offer that has business value such as profit or inventory reduction.

Targeting = Vendor Driven
Personalization = Consumer Driven

The Technology of Needs Inference

Now that we better understand the main goal of personalization, namely needs satisfaction, how do we accomplish it? What follows is a framework within which you can think about different need fulfillment approaches.

To start, let's consider the set of potentially observable behaviors and activities, or inputs, upon which we can base our inferences. Inputs related to users can be broken down into four categories:

1. **Behavioral signals** spanning present and past, implicit and explicit, online and offline activities
2. **Environmental context** such as location, device, time of day
3. **User attributes** like demographics.
4. **Social connections** such as Facebook friends

Many companies also invest in customer segmentation which is essentially a user attribute derived from one or more of the above inputs. A fifth category of biometric signals like heart rate, skin conductance, eye tracking and voice analysis has been explored in research settings for some time and applied to emotion and intent inference, but has yet to make its way into mainstream technologies. But, let's not forget the other side of the matchmaking game: the product. A personalization system not only needs to understand the users' needs, but also which needs are best filled by which products. There are three main categories of input related to products (or any other type of content):

1. **Product attributes** including anything objective about the product or content
2. **Metadata** which is expert-driven crowd-sourced descriptions and subjective categorizations of products
3. **Behavioral signals** such as how users interact with or consume products.

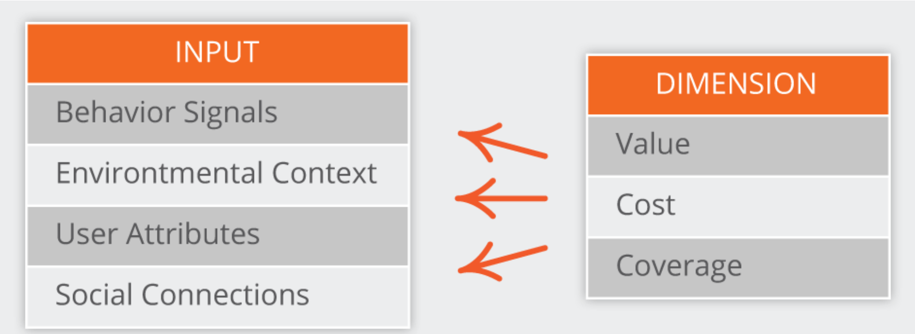

Figure 4

Even without biometrics, it's an immense amount of potential data. In general, any input should be considered along three key dimensions: value, coverage and cost.

Value is ultimately how reliably an input predicts a user's response to any one of many system outputs. However, the value of any one input must be considered in the context of other inputs. For example, knowing a user's gender may seem like a very valuable input for an apparel site. But, if a user's first behavior is to go to the men's clothing section, all of the sudden it stops mattering so much whether they are male or female; you want to show men's clothes, not women's. Data accuracy is another consideration that affects input value in that inaccurate data can be worse than no data.

Coverage relates to how often the input is available. If you decided gender was a valuable predictor of

> *"Mr. Siegel, a former marketing executive at Urban Outfitters experimented with gender personalization on the Urban Outfitters site. But it roundly backfired. It turned out that many female Urban Outfitters customers regularly bought men's items and they took offense at being subjected to gender-based marketing."*
> *– Slipstream, June 23, 2012*

behavior, yet you only have that information on a small proportion of users (e.g., most users are anonymous), it may have a low overall usefulness to a personalization system. As another example, technical challenges around user identification impact coverage and accuracy of past behavioral data.

Finally, **cost** tells us how much money or effort you have to give up in order to get at and process the data you want. Does the predictive value of that input offset the cost in setup, maintenance, computation, storage, monitoring and the potential demand placed on your user or your staff to provide it? Consider, for example, the cost of generating, updating and validating metadata - it might already be obvious at this point, but more is not always better when it comes to data.

Using the needs framework from earlier, let's drill down on the value dimension of inputs. What if, the user looking for the blue and green Grateful Dead T-shirt typed "blue and green Grateful Dead T-shirt" into your search box or had spent significant time looking at a bunch of different blue and green Grateful Dead T-shirts. Clearly these behavioral signals of current intent should be more valuable and trump anything you might know about the user's context, attributes or social connections. In this case, the user has come to you with a very specific, objective need that you want to respond to with relevant recommendations. Conversely, if the user came to your home page and gave you no clear behavioral indicator of current intent, then making intelligent guesses about what his needs might be based on past behaviors or user attributes is a reasonable strategy. While it's not completely black-and-white, a personalization system should do its best to infer the specificity of the user's needs and respond accordingly with focus on the right inputs and ultimately with the right amount of relevancy versus discovery.

The second half of the technology picture is outputs. After taking in and assessing all of the inputs, what actions and outputs should you take? I define outputs here as all of the various ways a personalization system can dynamically respond to a user. This includes recommended products, offers, informative and educational content in the form of documents and videos, messages, social validation (such as reviews), images or even direct requests to the user for information. Sometimes an output may trigger a whole new mode of interaction, such as a chat or email. Outputs can also be the look and feel or organization of a page and typically occur within a particular context or personalization zone, for instance on the homepage or within an email client. A user's response to any particular output then becomes an input to the system, and the behavioral feedback loop is created which is crucial to any personalization engine.

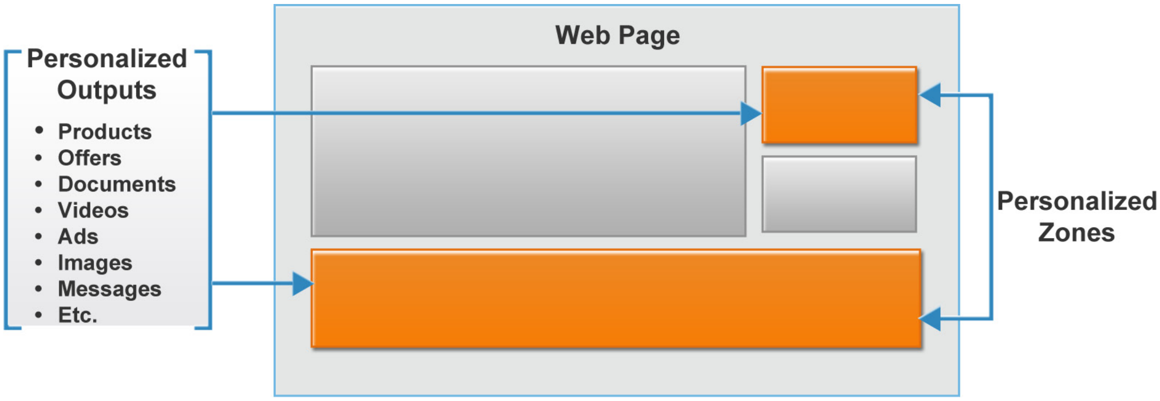

Personalization Output Zones

Figure 5

The Science and Art of Needs Inference

On their own, inputs and outputs are meaningless. The science and art of personalization lies in connecting the two. In other words, given everything we can observe about a user, what is the right action to take at this moment to best ensure satisfaction of his needs?

At a high level, there are really only two approaches: manual or automated. Manual approaches generally require the creation of rules that map a particular set of observed inputs to an output action. The advantage of manual approaches is that a human expert can directly encode her assumptions and knowledge into a human readable set of directives (e.g., "when you encounter a male age 30 – 40 from California who searches for T-shirts, recommend product X"). Multivariate or A/B testing is often layered on top of manual approaches so that different versions of rules or rule sets can be tested against one another for efficacy. The downside of these approaches is that the rule sets can grow arbitrarily large and complex leading to significant investment and challenges in maintaining and updating rules, including conflict resolution. There are also significant challenges in responding quickly in situations where either user needs or product catalogs change frequently or unpredictably.

To overcome these challenges with ever-growing user populations and products catalogs, the vast majority of modern personalization strategies include at least some degree of machine learning and automation. Here, the goal is for the system to figure out on its own how to connect inputs to outputs. Don't be fooled, though. This does not mean that humans have been removed from the equation when machines are asked to automate personalization.

First of all, rules don't go away completely. It's almost always the case that some rules need to be layered on top of automated systems to accommodate specific business needs or constraints. For example, forcing the output to display a mix of products, videos, or documents instead of those from a single category which can look boring even if they are the most relevant results, or telling the machine to boost the importance of new products so they get exposure sooner, or even blacklisting certain brands or products on competitor product pages. With better understanding of the mechanisms of personalization, one realizes that machine learning is both an art and a science. Beyond rules, there are three ways in which human decisions strongly influence the success of any machine learning strategy: choice of technique, inputs and features. There are myriad machine learning techniques, and you have probably heard of some of them — neural nets, logistic regression, collaborative filtering and Bayesian probabilistic modeling to name a few. Each of these uses a particular set of mathematical representations and learning algorithms. They approach the problem with a certain "mindset" that influences what can and can't be seen in the data. Machine learning is about finding the right generalizations from past data to effectively predict future data; each technique fundamentally generalizes in a different way.

> *"The ability of a machine learning algorithm to effectively predict and make recommendations based upon any combination is dependent on the constructs provided by humans."*

Think of these techniques like workers with specific skill sets. You want to choose the worker with the skill set that best matches the problem you are trying to solve. A worker with a sub-optimal skill set may require more time or direction to get the job done. There is no straight-forward answer for matching skill-set to problem other than to look at previous successes with similar problems. Yet, even a skilled worker without tools has difficulty. In the case of machine learning, the machine cannot decide which tools to use. Humans must first define the toolsets: inputs and features.

Inputs are the raw data, and the **features** are how you combine or organize those data. For example, a user's birthday is an input. However, if you give that to the machine learning algorithm directly, it still has to figure out if the day of the week matters most, and if people with similar birthdays act similarly, etc.

In general, inputs are combined to create features. In this case, birthday may be used to derive an age group feature, such as < 18 years old, 18-30, 30-50, 50+, etc. The machine is then given the much simpler task of figuring out if and how each category of age group responds. Another example of a feature is the combination of two or more inputs. For example, segmentation like "young urban professionals" is really a feature that combines multiple inputs such as age, demographic and location in a meaningful way.

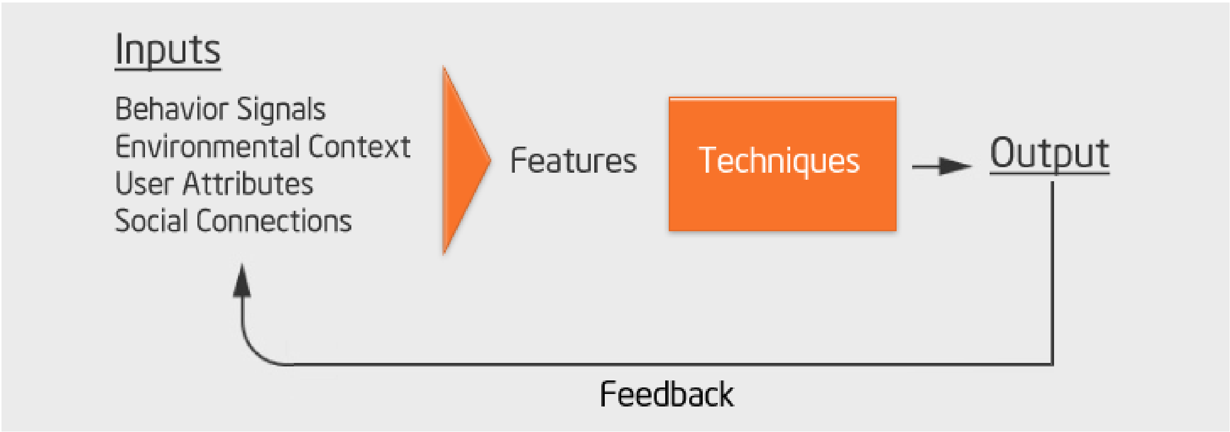

Figure 6

The ability of a machine learning algorithm to effectively predict and make recommendations based upon any combination is dependent on the constructs provided by humans. Give it great features, and it may be able to learn quickly. Give it bad features and it may take a long time or never finish. Give a skilled worker bad tools, and they have a hard time doing a good job.

This is important to note because many people think of machine learning as a super-powerful intelligence that can take any raw inputs and just "figure it out." Yes, it may be able to do something with the raw materials, but it is up to humans, for the most part, to craft the tools and features that will truly enable good machine learning.

BIG DATA AND THE MACHINE LEARNING CHALLENGE

Misconceptions of machine learning can run fairly deep. People not only have unrealistic expectations of what a machine can automatically figure out, but truly impossible ones. I sometimes think of this as the **omniscience vs. intelligence fallacy.** Intelligence is the ability to generalize and learn useful patterns from data and/or experiences. This is what we humans do naturally every day. But intelligence requires data. Very often the data that the machine learning system has access to is insufficient for anyone, even a super sophisticated intelligence, to generalize from. This is known as the problem of **data sparsity**. Unfortunately, in the real world data is more complex and more diverse across users. Each person that comes to your website browses a set of pages and has a certain set of profile attributes. Each user and what they choose to look at and buy is different from one user to the next. Think of it like a really big table where each person has thousands of possible cells and only 20 filled in and the 20 that are filled in are pretty different for everyone. Yes, you have tons of data, but the data are mostly different. If you plotted this data, it would have a lot of white space. The right machine learning techniques can get pretty good at spotting patterns in large volumes of sparse data but that doesn't mean that the machine can see something that is not there.

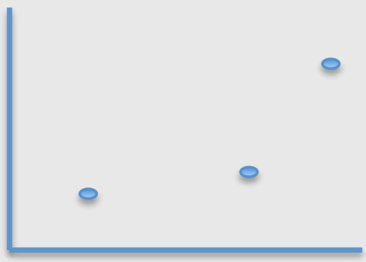

Data Sparsity

Figure 7

Another common mistake is what's known as **over-fitting**. To explain this problem, I need to take you back to high school algebra for a brief minute. You may remember that any two points on a graph can be connected with a line. You may also remember that any number of points scattered across a graph can be perfectly fit by a polynomial curve if it's allowed to be arbitrarily complex. In machine learning, the goal is to find the simplest curve that fits the current data points, but more importantly predicts where new as-yet-unseen data points will fall. If the machine chooses a curve that fits all of the current data points well (usually by employing an overly sophiscated model) but then does a bad job of predicting new data, it has succumbed to over-fitting.

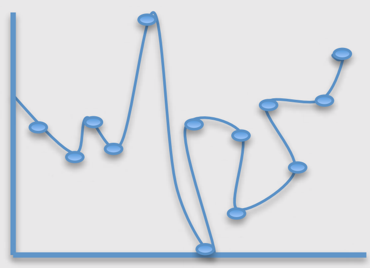

Over Fitting

Figure 8

I can't tell you how many times I have heard people run through scenarios in their head like the following: "So you mean if I knew that Fred had searched flowers and that his mother's birthday is next week, then I could predict that he's looking for a gift for his mom?" Probably not. The problem is that the model you have constructed in your head fits perfectly with the one or two examples that you have in mind, but when you apply that same generalization to all of the other examples, you're not thinking of, it fails more often than it works. You have succumbed to over-fitting. Strive for intelligent generalization in your personalization strategy, not omniscience.

Decision and Influence

Up to this point we've used the framework of need satisfaction and match-making to think through personalization. However, this isn't the full picture. The reality is that inferring the user's need and pinpointing the perfect product or set of products to meet that need is not always enough. A customer may not know that perfect fit when they see it. He may need additional supporting information on the product to truly understand how well it will meet their needs. In short, he needs to be *convinced*. He may also need to feel confident that he has fully considered the alternatives and may require some form of social validation on his choice. For this reason, it is important to think through what information you can provide a customer in support of their decision process. Comparison guides, testimonials, reviews, information sheets and videos are useful decision support tools. Triggering new modes of interaction, such as chat, at the appropriate time may also suit the purpose. Acknowledging the decision process can even extend post purchase, such as providing customers with follow-up validation and information which can help to reduce buyer's remorse.

Personalization strategies — whether manual or automated — often jump too quickly to the conclusion (i.e., the purchase). Think about a good salesperson. Even if she knows that everyone ends up buying the Nikon N367, simply presenting you with that product and effectively saying "here's your answer - buy it!" won't necessarily help. Instead, she will likely walk you through a few alternative cameras, each with different advantages/disadvantages and price points that could meet your needs.

A common mistake in thinking about the technology of personalization is to approach personalization as a prediction problem. The reality is that a machine which is perfectly able to predict a user's next action or purchase would have zero value. For example, if I'm 100% certain that you're going to buy Harry Potter, then predicting that you're going to do so and recommending Harry Potter to you has no real impact — you're going to buy it either way.

> *"The real game is influence, not prediction."*

The real game is influence not prediction. It is actually a much harder problem to solve from a machine learning perspective. Rather than simply predicting what a user's next action is going to be, the machine needs to predict what a user's response would be to every possible action it could take, and then to choose the action that changes user behavior in a way that optimally satisfies both the end-user and the provider.

The flip-side of decision is persuasion. The word persuasion may get a bad rap, but it is the reality of any commerce transaction. You want to present products that you want to sell now in a way that will lead to purchase conversion. Sometimes you even need to proactively reach out to customers and retrigger attention to a particular need, such as through email or ad retargeting.

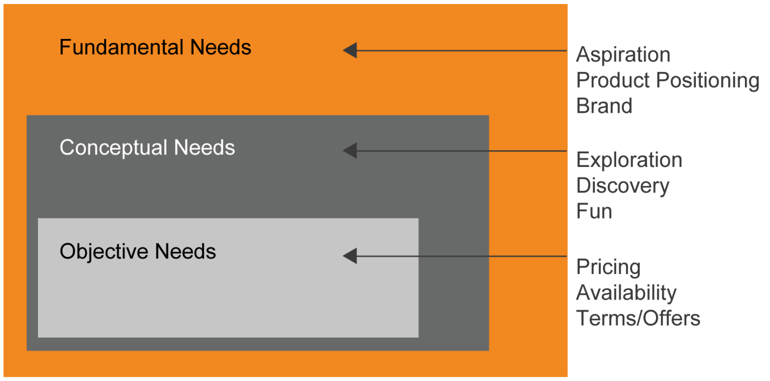

Figure 9

Shoppers respond differently to different messages and persuasion strategies. Some are more motivated by price, while others are more motivated by social reinforcement or perceived scarcity. This is sometimes referred to as a **persuasion profile**[5]. Messages must be tailored according to the stages of the buyflow. For example, a real or perceived discount may be best placed in the middle of the buy flow, after some consideration has been given, but before a product is chosen. This is hard to achieve, but personalizing the right persuasion strategy for the right person at the right time can make a big difference in conversion success.

THE FALLICY OF ONE TO ONE PERSONALIZATION

Connecting inputs and outputs is all about creating the right generalizations — either manually or automatically. The good news is that there are excellent generalizations to be made. As much as we think about personalization as one-to-one, the truth is that we are actually more the same in our needs and interests than we are different. For example, I may like science fiction movies and you may like historical dramas, but compare me to other sci-fi lovers, and I look very similar.

Going one step further, I like heady sci-fi with low violence, but there will almost always be a large reference group of people with whom I look more alike than different. For any particular need, interest or preference that I have, there is almost certainly a reference group that can be tapped for a better understanding of my need and how to fulfill it. In fact, it is my combination of interests that makes me unique, not my interests themselves. This is a good thing for personalization technology, because there is no way to learn if we can't generalize.

In everyday life, we don't treat every new person we meet as a completely unknown entity. Our brains are built to stereotype and generalize. Why? Because it is the only way we can function. Behaviorally, we apply a basic set of rules and categorizations learned over a lifetime of interactions with people to quickly figure out which categories or subcategories a person falls into and we act accordingly. That gets most of us about 99 percent of the way there. The more understanding we develop from our aggregation of experiences, the better able we are to understand any one individual. You may be thinking

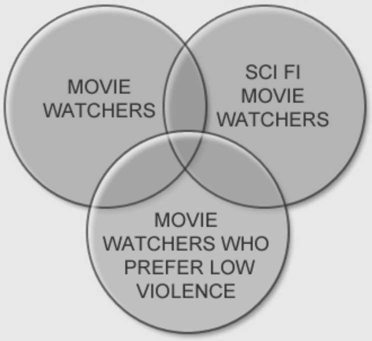

Figure 10

that you are not exactly like my stereotype, and this is true. Individuals use this categorization by first applying the stereotypes. Next, we learn to look for and apply only the minimal ways in which an individual differs from the stereotype. Within personalization systems, this is accomplished through a feedback loop. First, each individual is treated the same as anyone else who meets the unique combination of interests. Then, we correct ourselves if users don't respond as predicted. Sometimes that means reassigning a user to a different category because we may have guessed incorrectly from the start. Or, more likely, it means we have to watch and observe in order to learn additional unique characteristics of that user in the present moment.

Building Relationships through Personalization

Let's go back to where we started: the concept of uniqueness. Why do we intuitively feel the importance of treating each individual as unique? Social science gives us some answers. As humans, we have a fundamental need to be acknowledged and understood as a valued member of our community that transcends any particular utilitarian need. Think of the barista who knows our name and favorite drink. Any salesperson who takes the time to truly understand our needs and respond accordingly holds us in enough "positive regard" to invest the energy in doing so. And it makes us feel good even when we recognize that we are being manipulated.

Now here's the trick: psychologically, we don't distinguish between being recognized by a real person or a "social actor," such as a robot or even a website. A wide body of research in this area has shown again and again that we subconsciously apply the same rules and biases to interactions with machines as we do to interactions with other humans (Reeves & Nass Media Equation[6]). That's why we feel good when a site greets us with, "Hello Scott," even though we understand logically that we are not being addressed by a person. It's why we feel good when a site presents products and information specifically for us. We want to feel respected and recognized as individuals, even if by machines.

Whether we acknowledge it or not, user experiences with websites are social interactions. Our websites are our salespeople, and the salespeople can cultivate positive or negative relationships with our users. And, as with any relationship, trust is at the core. Building trust in interpersonal relationships requires both an understanding of one another's needs as well as interest and action in meeting those needs. A site which responds to our needs understands us, and feeling understood is known to have positive impact on well-being[7]. This lens can also help us understand the impact of doing a bad job of personalizing. When a salesperson or website provides us with recommendations that are completely off-base, we are forced to draw one of two conclusions: either the social actor is incompetent, or it doesn't care enough about us to do a good job. Neither is good, but like many things in psychology, lines become blurred. The inherent ambiguity here means that we generally respond to poor recommendations with both distrust and frustration at its general capabilities. Ultimately, this can lead to abandoning an interaction or more broadly, a relationship.

For machine learning, however, getting it wrong can actually be a very positive thing. Part of learning is making mistakes. There's no way around this, unless you already know the answer and tell the machine exactly what to do in every case. Sometimes, the quickest way to learn is to have the freedom to make lots of mistakes - the freedom to fail. But as an ecommerce site, bad recommendations and bad personalization can negatively impact your users and send them bouncing to a competitor's site. In the machine learning literature, this is known as the **exploration/ exploitation tradeoff** [8]**.**

> *"Bad recommendations and bad personalization can negatively impact your users and send them bouncing to a competitor's site."*

There is always a decision to be made: do I exploit the knowledge I've learned so far and make the best of it even if there's more I could learn? Or do I try out new things to see if I can find something better? This tradeoff is even more salient when context changes, such as when new products or offers are introduced. The more conservative you are in allowing the machine to try new ideas, the longer it will take for the machine to learn and optimize. If you are in a dynamic environment and play conservatively, the machine may remain so far behind the curve that it never does a good job of prediction or categorization.

One final personalization consideration is privacy. Beyond regulations and legal issues surrounding personally identifiable information (PII) there are real impacts of capturing and leveraging data that need to be carefully considered. Remember when we discussed inputs and costs? Sometimes the biggest cost can be to our customer relationships when we try to obtain an input or misuse consumer data in outputs. If we are building a long-term relationship with customers and having an ongoing interactive dialog, then we need to look beyond the transaction to how using data might impact the relationship we are working so hard to develop.

As in our social interactions in the real world, we earn the right to more intimate, personal relationships with our customers by building trust. Without trust, access to users' PII and other personal data can lead to uncomfortable feelings of vulnerability and invasion of privacy. Trust is created and managed by the responsible collection and use of personal information. Trust, however, can quickly be destroyed when this personal information is used or exposed in inappropriate ways. For example, knowing the consumer is over 65 and then overtly showing them products or offers for "seniors." If the consumer did not openly share their age with you, this could shake the trust you have built and cause them to abandon their session or even their relationship.

On the other hand, since a shopper's context changes frequently, it is possible to gather many other inputs that inform personalization more so than PII. For example, let's go back to the example of shopping for dad's coffee maker. It does not really matter what the personally identifiable attributes are. What matters is that the current context is understood and the right types of coffee makers are presented as choices. Similarly, stories abound of consumers being personally profiled by a technology, lets look at the 2002 TiVO example[9]. In this noted case, the TiVO wrongly assumed the sexual orientation of a particular subscriber, much to their dismay and ultimately the dismay of TiVO which was mocked for the issue.

We've talked a lot about our users' view on the relationship, now let's examine yours. True, your short-term need may be to close the most profitable transaction possible, but often times converting a customer and having them feel good about it is more valuable in the long term. Personalization is not just about conversion it's about building customer long-term value and loyalty.

Conclusion

We've covered a lot of ground in this paper. If it wasn't obvious from the beginning, it probably is now: personalization is a complex concept and process. But, if you get lost, it is always good to remember that personalization is fundamentally a human endeavor, both when it comes to user experience as well as the technology and science behind it. Personalization systems are not magic – and should not be expected to be magic. They are an encoding of human-conceived strategies for connecting inputs to outputs. The game is not simply about throwing more data into the mix; it's about being smart with how you select and organize inputs, the set of possible outputs and actions you enable, and how you frame the problem to the machine so that it can analyze that data.

As Albert Einstein once said "Everything should be made as simple as possible, but not simpler." Nowhere is this more true than in the world of personalization. Your goal is to make the process as simple as possible for your customer, even if the way you do that, is not very simple at all.

What It All Means

When thinking about approaches to personalization, consider these pivotal issues and they will steer your strategy in the direction that makes sense for both you and your customer.

1. Do the best you can to understand your user's needs and satisfy those needs with your products.
2. Help the user through their process; don't try to short-circuit the buy-flow by attempting to be "too smart".
3. Never lose sight of the fact that you are building relationships with your users. Treat them with respect. Focus on them and don't overstep the bounds of the trust you have cultivated.
4. Be thoughtful about your personalization strategies. It's not about throwing a bunch of data into a machine and pressing "go".
5. Take the long-term view. Personalization is not just about conversion, it's about building customer long-term value and loyalty.

End Notes

1. For more on Maslow's Hierarchy of Needs, see Abraham Maslow's paper, "A Theory of Human Motivation"; 1943. Originally published in Psychological Review, 50, 370-396.
2. For more on James McQuivey's 2008 speech on "Satisfying Consumers for the Next Decade." See Patricia Seybold's blog post at http://outsideinnovation.blogs.com/pseybold/2008/10/highlights-from-james-mcquivey-forrester-research-speech-on-satisfying-consumers-for-the-next-decade.html.
3. Although conceptions of personalization tend to focus on the need for uniqueness and individuality, connection and conformity is often a stronger and more reliable need. As such, the best way to meet a particular individual's overall needs may actually be to provide them with the products that other people like, along with social proof to evoke the notion that a purchase will enable them to fit in. From a needs-based perspective, conformity is not the antithesis of personalization but rather an aspect of it. See J. Wang and J. Lin. Are personalization systems really personal? Effects of conformity in reducing information overload. In Procs. of the 36th Hawaii Int. Conf. on Systems Sciences (HICSS03), pages 222-

222, Hawaii, USA, 2002.

4. For more on James McQuivey's 2008 speech on "Satisfying Consumers for the Next Decade." See Patricia Seybold's blog post at http://outsideinnovation.blogs.com/pseybold/2008/10/highlights-from-james-mcquivey-forrester-research-speech-on-satisfying-consumers-for-the-next-decade.html.
5. For more on persuasion, you can look at M. Kaptein and & D. Eckles 2010 work called Selecting effective means to any end: Futures and ethics of persuasion profiling. Proceedings of Persuasive Technology 2010, Lecture Notes in Computer Science., published by Springer.
6. Byron Reeves and Clifford Nass propose that we respond to communication media as if they were human. That people respond to media using the same rules that govern face-to-face interpersonal interactions with other people. Read more about their work at http://en.wikipedia.org/wiki/The_Media_Equation. Or you can review their book Reeves, B., & Nass, C. (1996). The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places by Cambridge University Press.
7. There is a variety of research that has been done on the topic of being understood and its impact on perceived well-being. See the research abstract by Janetta Lun, Selin Kesebir & Shige Oishi from the University of Virginia. Their two-week online diary study (N = 135) examined how feelings of being understood and misunderstood are related to well-being and health symptoms. Results showed that people who reported feeling more understood and less misunderstood in their social interaction also reported greater life satisfaction and fewer physical health symptoms on the same day and the following day. http://people.virginia.edu/~sk8dm/Presentations/Lun%20et%20al-2007-APS.pdf
8. There are numerous scholarly articles written about the exploration/exploration trade-off. If interested please reder to Exploration- exploration trade-off using variance estimates in multi armed bandits by Jean-Yves Audibert, Rémi Munos, Csaba Szepesvári as well as Near- Optimal Reinforcement learning in Polynomial Time dated 2002, by Micheal Kearns and Satinder Singh published by Kluwer Academic Publishers.
9. In November of 2002, The Wall Street Journal famously published an article called "If TiVO Thinks You Are Gay, Here's How to Set It Straight". In the article they discussed how TiVO personalization incorrectly profiled a number of users. You can read an excerpt here, but registration is required to see the whole article. http://online.wsj.com/article_email/SB1038261936872356908.html

This page is intentionally left blank.

Baynote

PERSONALIZED SHOPPING EXPERIENCES

Headquaters

BAYNOTE INC.
333 West San Carlos #700
San Jose, CA 95110 USA

USA Sales & Support
866.921.0919 x1 Sales
sales@baynote.com

877.803.2165 Support
support@baynote.com

1.866.921.0919 toll free

1.408.213.9601 interntional

1.408.973.9484 fax

www.baynote.com

About Baynote

Baynote is a leading provider of personalized customer experience solutions for multi-channel retailers. Using Baynote's patented approach, retailers personalize the shopping experience in the moment across touch points, increasing consumer engagement, conversions and order values. Based in San Jose, Calif., with offices in the U.K. and Germany, Baynote's personalization solutions are trusted by more than 300 of the world's most well-known brands, including Anthropologie, Bluefly, Campbell's, Crate & Barrel, Dell, J. Crew, Jockey and Urban Outfitters.

For more information about Baynote, visit http://www.baynote.com; read our blog and follow us on Twitter @Baynote.