Know the Room

*Learning your Internet audience **before** you spend marketing dollars*

By Ian Lurie, Portent Interactive

December 20, 2002

One piece missing from the Conversation Marketing method I'm so fond of is what to do **before** you enter a conversation. Assuming you have limited time and resources, how do you maximize the chances that the conversations you enter will be beneficial? Would you rather talk to 2 people and get two leads, or talk to 1000 people and get two leads? The answer, I think, is obvious.

What is Conversation Marketing?
If you haven't already, you may want to read Conversation Marketing before you go on. It's a short article, I promise, and it places this article in context. You can find it in the Library section of Portentinteractive.com.

On the Internet the problem is even more serious. You've got millions of potential visitors out there, and of the thousands of visitors you may get you have a smaller fraction of real potential customers. How do you assure that your web marketing initiatives will reach the right people? And once you do reach them, how do you make sure that you're sending the right message? Somehow, you want to separate good potential audiences from bad ones, and use that information to best target your communications campaigns.

In the world of conversation marketing, it pays to know the room, first.

That means you need to build a profile for your best audience, and apply that profile to your conversation marketing efforts. Doing so has three huge benefits:

Generate the maximum response with the minimum effort

If you know the room, then you can figure out the best ration of contacts to conversions. A **contact** is one piece of collateral – an e-mail, a banner impression, a piece of targeted content on a web site –

sent to one person or group. Depending on your goals, a **conversion** happens when the person you've contacted takes the desired action – buys a product, downloads a paper, or just views your web site.

Let's try an example: Say you have a database of 10,000 customers. They've all bought something from you in the past. Your company is about to send out a special offer via e-mail. Now, you can send the e-mail to all 10,000 customers, and get a response rate of perhaps 5%, or 500 conversions. But there are costs associated with an e-mail blast to your entire customer base:

> The angst factor: Some customers will get angry at the e-mail, even if they opted to receive it. These customers will be less likely to buy in the future.

> Bandwidth: Sending out an e-mail costs money. You have to pay for the bits, for the server time, etc.. Even if you own your own e-mail server and have lots of bandwidth, you're going to slow that server down and use bandwidth that might be better applied elsewhere in your organization.

> The frustration factor: Chances are you have a boss. Or a client. **You** may understand that getting 500 responses out of 10,000 people is pretty good, but will **they**?

If, on the other hand, you had a profile – a model – of the ideal recipient for this special offer, you might be able to narrow the list of potential customers to 1,000, and **still** get 500 responses. The recipients are more likely to appreciate the offer, you use fewer resources, and you get to tell the boss you just got a 50% response rate. Know the room, and you can attain marketing bliss.

Reduce Marketing Costs

Apply the same hypothetical to a banner campaign, or paid keyword sponsorship on a site like Overture, and you start saving hard dollars. If you can get the same response from 10000 banner impressions on two web sites as from 100,000 on six sites, why not do it?

With keyword sponsorships the same rule holds – if you're paying by the click or the view, wouldn't it be great to pay only for keywords that will generate a response, and only bid on the minimum position necessary to get a response (1st on the page, 2nd on the page, and so on)?

If you know the room, you can target your communications efforts, and minimize costs. You can spend the money you saved on Internet consultants. Ahem.

Increase Customer Satisfaction

I alluded to this above. Which do you prefer, as a consumer? Checking your mail or e-mail and finding 10 random marketing promotions that don't interest you? Or finding one promotion that's really compelling?

Take me for example – I'm a cyclist. If I get a flyer in my mailbox telling me about a 20% off sale at my local bike shop, I'm a happy guy. I even **appreciate** the notice, because I might not drive by the shop every week, and I'd otherwise miss the sale. The countless ads I receive for cheaper antacids (I have a favorite brand, thanks), dog food (I have 2 cats), and power tools (I'm totally helpless with hardware), on the other hand, just make my head hurt.

If you send the right message to the right person, they will always appreciate it, and you build loyalty as a result. Send the right message to the wrong person, though, and the best case is that they ignore you. The worst case is that they ignore from now on, even if, next time, you offer something they might want. Know the room, and people like you better.

Convinced? I hope so. But how do you actually **do** all of this?

Methods: How to See the Room Before You Arrive

To know the room on the Internet, you have to build a profile of the ideal target audience. There are a few basic methods to do this. I'll discuss each one briefly and point out the advantages and disadvantages of each:

> Your gut
>
> Data query
>
> Data mining

Whichever one you choose, you need to start with a **hypothesis** – an assumption that includes what you want to know, and the information you can use to learn it. Typical hypotheses include:

> Given what I know about my customers, one neighborhood might really want my newest product.
>
> Given my sales database for the past year, some set of fields or attributes defines the ideal target for this campaign.
>
> Given weather reports for the past year, combined with lift ticket sales, I can predict which forecasts lead to the highest number of tickets sold.
>
> Given traffic on my web site, search engine campaign records, and sales data, I can determine which keyword combinations will generate the most sales.

Prediction Versus Measurement

Of course, you can records statistics and measure results for these, but the goal here is to **predict**, not **measure**. Measurement is critically important because it determines whether your campaign worked or not, and can serve as the foundation for future analysis. But prediction lays the foundation of each campaign.

Once you have a hypothesis, you look for a set of attributes that define the best audience. Attribute sets that we've found on different projects include:

> Men and women between 20 and 30 years of age, who own a car.
>
> Women older than 40 who live on the East Coast.
>
> People who previously bought product A or B.
>
> Attendees at the last convention.
>
> Interviewees who stated they were afraid of heights.

Each of these attribute sets helped us create a profile of an ideal audience, and a more focused communications initiative.

Your Gut

You know your field, and can probably get a good feel for what works or what doesn't with your customers. You can probably judge, without even knowing exactly how, whether a given promotion will work on a specific group.

For example, if I owned a bicycle shop, I could probably guess that almost any house within 10 blocks of my shop is a good target for a marketing flier. Houses that are further away, though, are closer to my competitor than to me.

The gut-check method has two advantages:

It's cheap: You don't have to hire any consultants (hmph), buy any expensive data processing software or otherwise invest in infrastructure. All the tools are in your head.

It's flexible: Your brain is really good at recognizing strange, unpredictable patterns in non-numerical data. For this kind of 'soft' information, you may be the only judge of the room.

But the downside is serious:

It's inefficient: Going by intuition is OK if your potential audience is small, and your data set is simple. But if you have hundreds, thousands or millions of customers and/or many possible attributes, you may start to lose track, or give up all together.

It's imprecise: Once the attribute set gets large enough, you will probably unconsciously edit the set (psychiatrists have written tomes about this).

It's biased: See imprecise, above. Your own personal feelings about each attribute will **always** affect your interpretation of those attributes. That may skew your analysis and lead to less useful results.

So, the gut-check method is usually the least favorable. But there are times when it's the only way. If your data is totally subjective, for example, nothing else will work. Say you have a survey with several full-text answers, where the interviewee can write whatever they like. You **might** be able to use a computer to grab key phrases and analyze accordingly, but I doubt it. You have to read the answers and draw your own conclusions.

So, intuition definitely has a place, but it's best saved for purely subjective data.

Data Query

If you run any type of organization these days, chances are you have databases full of information: Membership records, sales data, site traffic reports, test results, etc. are treasure troves of useful data.

A Database By Any Other Name…
By 'database' I don't simply mean database software, such as Microsoft Access. A database could be on paper, in a spreadsheet, a word processor, or a conventional database. For the purposes of this article, consider a 'database' an orderly information collection, such as customer names and phone numbers, sales data or survey records.

You can easily query that data and look for patterns that indicate a useful profile. Data query is a three-step process:

1. Convert the data to a useful format: If your data is on paper, enter it into spreadsheet or database software. If it's in a spreadsheet or database program already, make sure that all data is properly formatted (numbers are in number format, etc.) and stored in a way that enables your queries. If you don't understand this step, don't worry about it – you'll have to find a database programmer or spreadsheet guru, anyway, and they can take care of conversion.

2. Run your queries: Next, you build queries based on your hypothesis and run them against the data. That will yield a smaller data set.

3. Analyze the result: Finally, look at the results, and see if a pattern emerges that indicates a good profile.

Let's take my bicycle shop again – if I run a query against my sales records to find all customers who purchase from my shop more than once, I find that the result only includes mountain bikers. Hmmm. As a result, I buy banner advertising on a Seattle web site dedicated to mountain

bike trails in the region. To further target that audience, I link those banners to a special page on my web site that talks about our mountain bike specials, and I include a big picture of me fixing a flat on a dirt road in the middle of nowhere. I know that the 'room' is filled with mountain bikers, so my ad buy targets that audience, I dress appropriately by demonstrating I'm a mountain biker, too, and I say the right things by talking about mountain bike deals.

Dress Appropriately?
'Dress Appropriately' is another rule of Conversation Marketing. Basically, it means a web site should be designed appropriately for its audience – a web site created for a muffler manufacturer will not look the same as a site designed for a rock band.

Data query has many advantages:

It's a wonderful way to look for patterns you may expect.

It's still fairly inexpensive: You can do it using something as simple as Excel™.

It's quick: Once you prepare your data you can query it many different ways.

It's scientific: You're working, typically, with numeric data or a limited range of non-numeric choices, so results are pretty unequivocal.

It's sexy: If you have a boss, or a client, you can persuade them that a specific model is correct with tables, charts, graphs and who-knows-what. You can prove you really crunched the numbers. Right or wrong, that holds a lot of value in today's organization.

It also has some limitations:

Expertise required: You're going to have to hire someone who understands databases, or learn yourself.

It can't handle subjective data: See 'Gut Check', above.

It's limited: Data query can only predict using patterns **you** thought of. If I didn't remember to check for the type of products customers purchased, then I wouldn't have known that mountain bikers were my best repeat customers. In my example that seems like a no-brainer, but data query can miss subtle patterns that have a profound impact. If you use data query, consider your queries very, very carefully.

Consider data query an ideal middle ground if you have lots of numerical data and limited resources, or want to test the analytical waters before you jump into data mining.

Data Mining/Knowledge Discovery

Data mining – also called knowledge discovery – is one step beyond data query. Instead of looking for patterns based on an ideal model you try to predict, data mining looks for patterns you might not have thought of. Basically, the computer looks for patterns by learning your data set.

OK, Back Off…
Before all you data mining experts deluge me with angry e-mails, realize that I'm trying to sum this all up in 3 paragraphs or less. Data mining is a topic that fills many, many books – but for our purposes this definition will do.

I won't go into the details – it would take another 300 pages, at least. But you can use data mining to find patterns you never would have thought of.

Back to my bicycle shop for a moment. If I put my sales data into a data warehouse, and then analyze it using one of several data mining methods, I might find something surprising: My best customers have kids. Turns out the little tykes wreck and lose equipment at an amazing rate, so mom and dad are back for new helmets, gloves, tires, bicycles and who-knows-what every few months. And, parents with older children who ride mountain bikes are a gold mine.

And, I find a separate pattern at the same time: Parents with only one child spend only 25% of parents with two children. Guess it's easier to keep track of one set of training wheels than two or three or…

I didn't start my data mining project planning to find these patterns – my data mining software found them for me. As a result, I have two much more precise models, instead of one broader one. My communications tactics can be that much more efficient.

If you're trying to know the room, data mining means you arrive in a limousine. It's the most precise, effective way to take vast amounts of information and build a strategy. The advantages include:

> Effectiveness: Because you don't have to guess at patterns ahead of time, data mining delivers the best results with numeric and limited non-numeric data.

> Efficiency: Once you prepare your data you can apply many different models and methods to it. You minimize repetitions.

> Completeness: Data mining will capture more potential profiles and patterns than any other method.

However, there are problems:

> Price: Data mining costs a **lot**. The software tends to be expensive (although you can use spreadsheets for some forms of analysis), and you'll likely have to hire a consultant.

> Complexity: Data mining is not simple. It takes a lot of preparation and careful review of results.

But, if you can afford it, data mining is absolutely the best option for knowing the room.

Build a Profile: Get Results

No matter which method you use, it always pays to know the room before you start a communications campaign. Most marketers use a little bit of all of the methods I've described in this article. But regardless of how you do it, if you build a profile of your target Internet audience, you'll know who your banner ads, search engine campaign, web site or other initiative is talking to. If you know that, you can increase customer loyalty, build conversion rates and make the best use of your resources.