


THE MARKETER'S GUIDE TO CRM AUDIENCE BUILDING:

What We've Been Doing Wrong All Along, and Why Plug & Play Deep Learning is the Solution



Digital Marketing Director
Tinyclues, New York



If your team sends targeted CRM campaigns and does NOT leverage Deep Learning to do so, then you know what I'm talking about... That sneaking suspicion that every audience is arbitrary.

In this illustrated guide, I'll show you the path to CRM greatness. Buckle up!

Have you ever felt frustrated with your team's CRM audience building process? Maybe you've wondered if your query criteria are too loosey-goosey? Or how the process could possibly be so cumbersome?. Or maybe you've simply found yourself asking, *In 2021, is there really no better way?*

Alternatively, are you only sending Batch & Blast newsletters, but considering implementing a targeted approach in the near future - and possibly feeling overwhelmed by the implications in terms of data, segmentation, and bandwidth?

Or maybe you're just curious to learn about AI applications in the field of B2C CRM marketing?

Whatever the case, I've got you covered. In this illustrated guide, I'll go through a step-by-step explanation of what's wrong with traditional audience-building methods, and a demonstration of why Deep Learning is the only way to do CRM Audience Building right (as well as a few other handy CRM marketing applications!).

But first, a quick note: this article is a little long. Here's what we're going to cover:

- A refresher on the typology of CRM communications
- The whole premise of targeted campaigns rests on them being as relevant as possible
- But traditional audience-building methods fail to achieve relevancy
- Deep Learning is the solution to achieve true CRM relevancy, because it can most accurately predict propensity to buy
- But readers beware: not all Artificial Intelligence is created equal
- So what does it take to achieve truly great results?
- If Deep Learning can predict propensity to buy, what are some real-world applications?
- Where to learn more on this topic

A refresher on the typology of CRM communications

To start things off, let's get some definitions straight, seeing as different organizations tend to have varying terms for different types of CRM communications.

CRM communications usually fall under one of these categories:

- **Trigger messages:** Communications that are automatically triggered by a customer's actions. Typical example: Abandoned cart.
- **Journeys:** Automated sequences of communications that are tied to the customer lifecycle. Typical example: New customer welcome series
- **Campaigns:** One-off messages that are not sent automatically, and that are either:
 - Sent to your whole customer list ("Batch & Blast")
 - Sent to a fraction of your customer list ("Targeted campaigns").

Today's focus will be on the latter. Campaigns usually make up about 90% of the CRM messages sent by a brand. They are used to communicate about new collections or products, sales, specific product categories, and more...

Typically, CRM teams start by implementing "Batch & Blast" campaigns - campaigns that they send to their whole customer base. However, they usually see the limits of Batch & Blast very quickly:

- Batch & Blast campaigns are not very relevant to many if not most of your customers, since you're sending the same content to everyone; therefore...
- You can only send a limited number of these campaigns each week if you don't want to overwhelm your customers with irrelevant messages and create a negative experience

The logical next step is therefore to move on to targeted campaigns, featuring

specific topics or offers to be sent to the subset of your customer base most interested in that particular content.

The whole premise of targeted campaigns rests on them being as relevant as possible

Why do CRM teams implement targeted campaigns?

The reasons are many:

- Targeted campaigns will, in theory, be more relevant and therefore provide a better customer experience with each message...
- ... which will positively impact engagement, conversion rates, and revenue
- ... and will decrease unsubscribes.
- Targeted campaigns also allow teams to communicate about more offers and products, therefore giving more CRM coverage to their whole product or offer catalog and make more stakeholders happy (think: category managers asking for their products to be featured in newsletters, partner brands requesting access to a retailers' customer list...). This is true not only because you can send a wider variety of messages (to fewer customers each); but also because, *if the overall relevance of the messages that you send is higher, you can send more messages to each customer.*

Attaining these objectives with targeted campaigns, however, is obviously subject to making the theory a reality: these campaigns do need to be as relevant as possible.

Success hinges on the assumption that the CRM team is able to identify, for each targeted campaign, the subset of customers that will be most interested in the campaign.

But traditional audience-building methods fail to achieve relevancy

So how does one make a campaign relevant? The CRM team needs to identify, for each targeted campaign, the subset of customers that will be interested in the campaign topic(s), engage with the campaign, and, ultimately, convert.

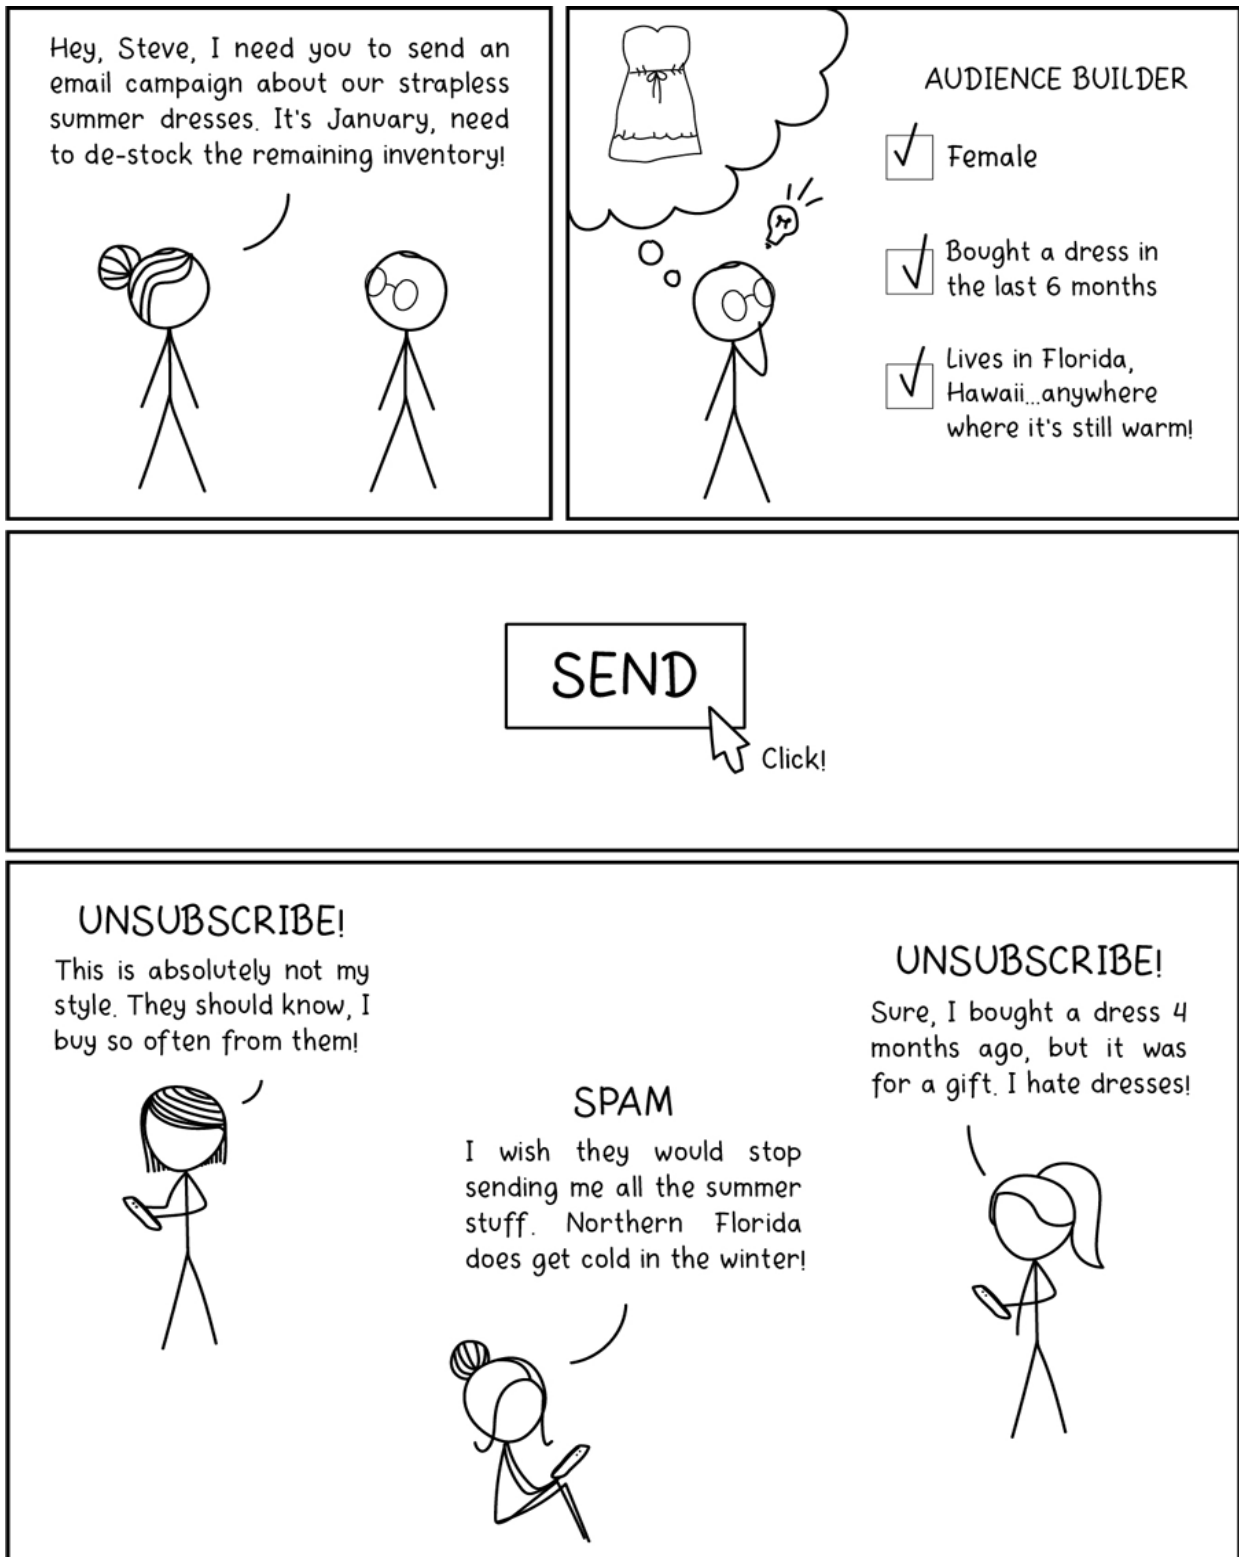
If this sounds easy, think again.

As marketers, we've all been taught to segment our market and put our customers in neat little boxes. We've practiced it so much, we forget that these customers are actual people, in all their complexity, with their lovable quirks, whims, ideas, spontaneity - in short, everything that makes them human.

Or maybe we don't forget... but what other options do we have?

With traditional audience-building, human bias is built into the methodology

In practice, this is what it looks like:



Of course, this is slightly caricatural. Hopefully the campaign above would generate some sales, and not just unsubscribes. But you get the point. The traditional, rule-based approach to campaign targeting has many flaws.

Bias is built into the methodology - twice!

Your CRM first-party data is super rich. You have so many data points about your customers, their behavior, and your products.

And yet when building audiences as described above, there's a double layer of biased simplification at play:

1. **You need to decide what dimensions are good predictors of purchase behavior.** In the example above, Steve decided that gender, recent purchase history, and location are the most important dimensions. But why? How can we be sure they're better dimensions to look at than age, first name, and browsing history? Or any number of other data points?
2. **For each of these dimensions, you need to decide what values to select.** In the example above: female, recent purchase of a dress, specific states. Once again, why? Maybe the most interesting feature of these dresses is that they are strapless. Wouldn't selecting recent buyers of strapless tops be a better choice? Maybe, maybe not. Who knows? For both dimensions and values, the selections will be guided by intuition (best case) or guesses (worst case), and, invariably, will be plagued by our assumptions and biases.

Therefore, you can guess my next point...

Performance will be mediocre at best

Because this method is only a very rough attempt at "predicting" relevancy of the campaign

- Some people are going to be *included* in the audience even though the campaign is entirely irrelevant to them.
- Some people are going to be *excluded* from the audience even though it would have been 100% relevant to them.
- This approach is rarely sophisticated enough to take *seasonality* into account, or *product lifecycle*. People who pre-order a new video game have a very different profile than people who buy the game months after it's been released as a Christmas gift. How would any traditional audience-building method ever account for this?

You're siloing customers instead of promoting cross-sell

Extrapolating purchase history to predict future purchases will unfortunately trap your customers into their starting categories and you'll be missing out on the opportunity to introduce them to additional offerings, even though those offerings might be relevant to them.

Reach can easily get ridiculous

One pitfall of building criteria lists in the hope of narrowing down your audience is that you're, you know, narrowing your audience. Actually, it's really common for marketers to realize that their final audience is too small. So then they have to compromise, add new segments, remove some criteria... ending up where they started: with a broad and undefined target.

You're not predicting propensity to buy NOW


For audience-building, predicting the propensity to buy a given product is actually not enough. You want to predict propensity to buy at the time of the campaign send.

It's not enough to say, "Steve has bought a dress in the past 6 months, therefore Steve will likely buy another dress someday" to include Steve in your dress campaign. Over a very long period of time, this statement stands a good chance of becoming true, but that doesn't really help the campaign that you need to send *tomorrow*.

It would be much more powerful to be able to say, "Steve has a very high propensity to buy a dress this week, so let's include him in tomorrow's dress campaign"

Alright, I hope we're now aligned on the limitations of traditional audience-building methods.

Luckily, there is a solution, and it is, of course, AI -more specifically, Deep Learning.





Deep Learning is the solution to achieve true CRM relevancy, because it can most accurately predict propensity to buy now

And when you think about it, it makes a lot of sense:

- CRM marketers have **a ton** of first-party data and it so happens that AI is really good at making sense of big datasets
- Our campaign targeting problem has a lot of similarities with one very popular use case for machine learning: image recognition.

What are the similarities of audience-building with image recognition you might ask? Well:

1. **In each of them we're trying to answer a simple question...** (*Is this image a cat? Is Steve going to buy a trip to Thailand this week?*)
2. **In both cases, we have very rich, complex data available to help answer the question:**
 - Images are composed of hundreds of thousands of pixels, each of which can take any of the 16,777,216 possible color values (in RGB).
 - CRM data... between customer attributes, product attributes, and transaction history, you easily have over a hundred columns in your table. Now just imagine the possible combinations... adding browsing and campaign history, tacking on temporality... your CRM tables have a depth of data that is sometimes hard to envision!
3. **We might think we can focus on a small subset of criteria to answer our questions...**
4. But we'll find that even if we spend a very long time honing a criteria list, **the success rate remains very weak.**



IS THIS IMAGE A CAT?


2 Ears

Furry

1 Tail

4 Legs

WHY THESE CRITERIA FAIL



IS THIS PERSON GOING TO BUY A DRESS?

Female

20-49 years old

Has bought a dress in the past 6 months

Has browsed dress website category

WHY THESE CRITERIA FAIL



5. However, in both cases, ***it is possible*** to generate a huge dataset of examples where the question has already been answered.
- **For image recognition:** there are datasets of images that either contain cats, and are labeled as such, or do not contain cats, and are labeled as such
 - **For CRM purchase behavior prediction:** the combination of your customer base, your transaction history, and your product attributes represents many examples where the question has been answered! "Did Peter buy a dress on December 3rd? "No" On December 4th? "No" Did Mary buy a dress on December 3rd? "Yes" On December 4th? "No" Did Jennifer..."

And we know that Deep Learning works amazingly well for image recognition. With a big enough dataset to learn from, the success rate is close to 100%. In some specific applications (medical imaging comes to mind), models are already more efficient than humans!

Deep Learning applied to image recognition (and in general) works much like the human brain does- the same way a child learns to recognize a cat not because he learns a list of criteria provided by its parents, but because they patiently show him pets and comment "this is a cat", "this is a dog".

Deep Learning applied to relational databases like a CRM dataset is much the same. It'll learn from the entirety of your database, the trove of data we mentioned earlier, without discarding any datapoint or dimension, without being burdened by arbitrary rules. It'll ingest all the edge cases that criteria would never capture, and incorporate them to tune ever-so-slightly its predictions.


Deep Learning will be able to detect subtleties like product lifecycle, seasonality, people's "style", people's "unpredictability"... without any of these dimensions being specifically inscribed in your database.

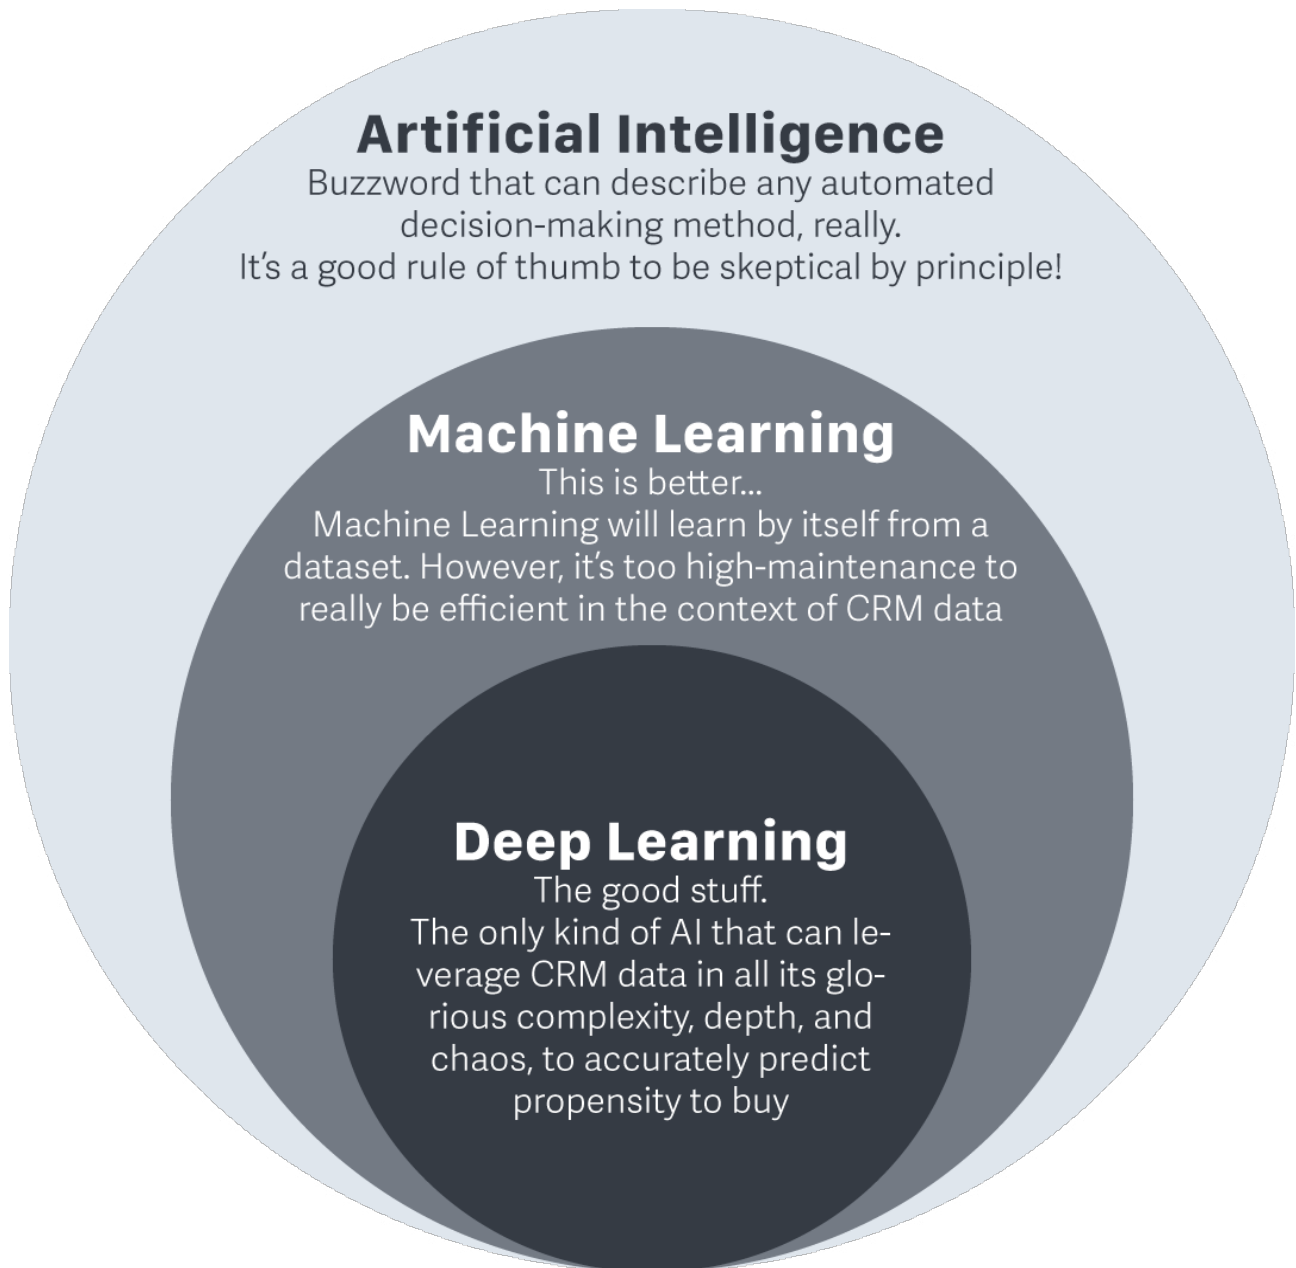
Deep Learning can accurately predict that Mary is going to buy a vintage floral dress, even though Mary never bought a dress from you before – or anything tagged as "floral" or "vintage," for that matter.

But readers beware: not all Artificial Intelligence is created equal

Of course, what I described above only applies to Machine Learning done right. You can't just upload CRM tables in an algorithm for image recognition in order to get purchase intent predictions.

And similarly, we all know that AI and Machine Learning are pretty much buzzwords at this point. So here are a few points it's good to have in mind when thinking about incorporating AI or Machine Learning into your CRM strategy (!):





- **Artificial Intelligence** is a very generic term that can include any automated decision making process, including rule-based decision trees (like the ones we described above, which we found to have so many shortcomings!). Interestingly, rule-based AI ("Expert Systems") was the dominant form of AI in the 1980s... leading to the second "AI winter" (dark age for AI when enthusiasm for AI decreased and investments slowed to a near halt).
- **Machine Learning** is a much more advanced field of AI. Research in Machine Learning started invigorating the AI field again after the second "AI Winter", in the late 2000s. No more hard-coded, rule-based decision making. Essentially, Machine Learning will learn its own predictive rules based on patterns it

sees in the data. As we've seen before, that's absolutely essential when it comes to CRM data and predicting propensity to buy! **However, Machine Learning still has an important drawback:** it is somewhat dependent on human intervention to make sure the data is structured and labeled in the proper way... Which, of course, can't be relied on in our CRM context.

- **Deep Learning** is an even more powerful subset of Machine Learning. Deep Learning essentially mimics the way our brains work. Deep Learning's "Artificial Neural Networks" can learn from any type of data, with minimal human guidance required to achieve outstanding results. These networks are built to be able to structure the data into meaningful representations and learn from them - no human intervention required!

Bingo! It sounds like Deep Learning is the appropriate form of AI for propensity-to-buy predictions.

Now, within the context of Deep Learning, important questions remain:

- What datasets are being used?
- How do the algorithms handle unclean, incomplete data?
- Within the datasets, what data do the algorithms learn from? Are data points from 5 years ago given the same weight as more recent data points? Why or why not? And how is this implemented?
- What do the algorithms optimize for?
- How resource-efficient is the technology? (Will it score your customer base in a day, or in a minute?)

As is so often true, the actual value is in the chosen application, and the execution.

That is why so many "Deep Learning" frameworks are actually open-sourced or commercially available. Even though they're really good building blocks, they will not get you very far on their own.

That is also a reason why, generally speaking, you should be suspicious of AI, Machine Learning, or Deep Learning applications that have not been developed for a very specific purpose, but for a wide variety of use cases. (Yes, I'm looking at you, Einstein and Watson).

Not all AI or Deep Learning is created equal.

So what does it take to achieve truly great results?

In order for Deep Learning to be applied successfully to CRM audience building, it needs to follow some basic principles, which we can think about in terms of:

1. **Efficiency:** the results need to be as good as possible; and
2. **Practicality:** the implementation needs to be as easy as possible

The one without the other would just mean that the application will never be used.

Let's get some of the obvious practicality questions out of the way:

- **The application should solve real life problems.**
- **The application should be adapted to a CRM team's workflow.**
- **The application should be easy to use by a CRM marketer, without the need for data science skills.**

These do sound obvious (they're why applications are called "solutions" and not "problems"!) but you might be surprised at the number of "solutions" that don't tick these boxes.

That being said, let's move on to the juicy, interesting parts.

The application should use all your available first-party data, in whatever state it is in NOW, even if it's messy, siloed, and all over the place.

- **Practicality:** Clean data does not exist. Most companies have started looking for a CDP, but it's a two-year project... Are we really going to wait that long? What if two years turn into three? And guess what we'll find out when the CDP is finally ready? The data is still not clean. Clean Data does not exist. We have to live with that, and achieve great results all the same. The good news is, that's perfectly achievable if the model has been developed with this constraint in mind.
- **Efficiency:** Data is important. Data is a signal. Even messy, unclean data. Let Deep Learning learn by itself what is a key signal and what is noise. It's what it does best, after all.

The application should not need rules or criteria.

- **Efficiency:** Rules would defeat the purpose of using a Deep Learning application. It would mean re-introducing bias into the predictions. Remember the "AI winter"...
- **Practicality:** Coming up with rules is time-consuming!

There should be no cold-start.

This one is easy, both from the efficiency and the practicality point of view. We want results from Day One, that's non-negotiable. The Deep Learning

model has to be able to learn from historical data and have great predictive quality from the start. (More about this in this article by one of our data scientists - Warning: it's a little bit more technical than what you're reading right now).

Seasonality, trends, and other changes in consumer behaviors should be reflected in the predictions.

It's mostly a matter of **efficiency**: If these trends and behaviors aren't reflected in the predictions, then the results are going to be mediocre at best. Purchase patterns are wildly different during the holiday season than, say, early spring! To best account for seasonality, the model should be tuned to pay special attention to (learn predictive clues from) the most recent purchases of the offer it's predicting propensity for!

It should optimize toward your preferred success metric - Revenue, in most cases, rather than open rates, click rates, etc.

It should be able to make accurate predictions even for customers for whom there is no fresh data.

In a given week, maybe 2% of your customers visit your website or make a transaction, giving you fresh data points. Over 3 or 6 months, maybe it's 20% or 30% of them? What about the others, the vast majority of your customers for whom you do not have fresh data? You still want to be able to send them relevant communications. Maybe they're even inactive, so the need for relevancy in order to reactivate them is even greater. That's the pitfall of many personalization / product recommendation solutions: they do OK for customers with very fresh intent data, and are absolutely counter-productive for the rest.

As a result of all of the above, the application should be able to make cross-category predictions, and go beyond straightforward purchase history extrapolation.

Because very few humans buy only dresses, or only rugs, or only flights to Miami. If John bought socks from you before, are you only ever going to send John emails about socks? If John visited your hotel in Chicago in the past, is he condemned to receiving "Chicago" communications forever? What if John loves to never visit the same place twice?

As you can see, there are many considerations that go into developing a Deep Learning model dedicated to predicting propensity to buy from a CRM dataset. So many, actually, that it's never worth it for a B2C brand to build it internally:

the level of expertise and development time required are prohibitive.

However, for teams of expert data-scientists whose sole focus for years is creating and refining such a model, it actually is possible... and it opens so many real-world marketing use cases!

If Deep Learning can predict propensity to buy, what are some real-world applications?

There are endless practical applications of being able to accurately predict propensity to buy for each customer in your base, for each product or offering in your catalog. Here are a few of them:

Audience Building

The first one, obviously, is audience-building (good thing, seeing as it's the topic of this article!).

Let's say your Deep Learning algorithm can rank all of your customers from most likely to least likely to buy a given product. Then building the audience could be as easy as:

1. Selecting the product(s) or offers showcased in your CRM Campaign
2. Selecting the desired volume of recipients OR taking Deep Learning's recommendation (sending to everybody who has a higher-than-average propensity to buy, for instance)

Marketing Fatigue Management

CRM teams send many campaigns each week, sometimes multiple per day. But they might have a rule that no customer should receive, say, more than 3 messages per week, or more than one message per day.

What if a customer has a really high propensity to buy each of the products featured in a given day's campaigns, though?

The propensity to buy information could be used by the solution to assign this customer to the one campaign for which his buying propensity is the highest, ensuring maximum relevancy while minimizing fatigue.

Multichannel CRM

If you can get propensity to buy by **customer x product**, there's no reason you couldn't get propensity to buy by **customer x product x channel**.

Therefore adding one more dimension to all of the above!

Identifying untapped “pockets” of propensity to buy

Maybe there's a specific subset of your customers who are really hot for a given product right now. It might not be visible in the KPIs available to the CRM team, because it's not yet among the “best sellers”... but there's a high demand concentrated among a small number of people. That's the perfect audience for a targeted CRM campaign!

This information could be used to create a “heatmap” of your product catalog, letting you visualize these opportunities easily.

CRM calendar planning

Should two products be featured in the same campaign, or in two distinct campaigns?

Well, the answer depends on whether the people interested in one, will also be interested in the other.

Guess how to get this information...

Where to learn more on this topic

If you've made it this far in the article, thank you! I hope you found it interesting and instructive.

Interested in learning more? Here's what you can do:

- Watch [this video](#) where the VP of CRM at Sally Beauty explains how he successfully deployed Machine Learning (*actually, it was really Deep Learning ;)) to scale a customer-centric strategy
- Read [this article](#) by Matt, our Product Marketing Director, for more on audience-building and targeting strategy
- [Send me an email](#) to set up some time to chat (you may have noticed I'm passionate about the topic and could go on and on ;))
- [Request a demo](#) of Tinyclues (because, as you may have guessed by now, every Deep Learning promise I described in this article... Well, our fantastic data science team has achieved it, and beyond).