Chapter 2.

w.

Artificial Intelligence Driven Design.

Brain food — Vol 4

By Joël van Bodegraven

awwwards.books

Chapter 2.

TRAINING AI WITH DESIGN

This second chapter was written by Pedro Marques & Joël van Bodegraven, Product Designers at Adyen.

Chapter 2.

- Introduction
- The importance of Machine Learning
- How does AI learn?
- Machine Learning, the first step to AI
- Machine Learning applied on design (or as a design pattern)
- The user on the loop
- How do you create this feedback loop in a ML mode?
- Design enabling AI
- Biases
- Join the ML playground!

Introduction.

In the first chapter, we've touched the surface of how AI affects our process as creatives and how important data is in delivering meaningful user experiences. In this chapter, we dive deeper into how we can train Artificial Intelligence (from now AI) with design.

AI relies heavily on data but how does it learn?

How does AI become intelligent?

The goal of this chapter is to introduce and get you started with several methods and techniques to make your AI smarter with design.

The importance of Machine Learning.

Code has always been that ‘invisible’ part of great experiences. With the dawn of an AI-driven society, Machine Learning has become a design decision as well. Machine Learning is a hot topic in our industry and changes the way we design and create our products. It offers tons of opportunities but also comes with some important challenges.

One of those challenges is that Machine Learning is now a UX-problem. By sharing how Artificial Intelligence learns and how you can have an impact on this as a designer, we hope to spark your imagination and boost your confidence so you can play a bigger role in your next AI-driven project.

How does AI learn?

There are several ways that Artificial Intelligence (AI) can learn and evolve. Generally you can say that an AI learns from (data) patterns, mistakes/errors and successes (trial and error). Opposed to conventional programming, Machine Learning starts at the end, so you don't code the solution, the machine itself learns from a set of rules.

Imagine you want to make the best burger in the world. Instead of trying to come up with your own recipe. You buy several great burgers and you start from there.

You bought a bunch of burgers and tell the machine what it is and you tell which parts you like the most. The machine takes it from there and deconstructs the burgers and comes up with the perfect burger. The 'magic' around Artificial Intelligence is, is that you might not even know HOW it is done.

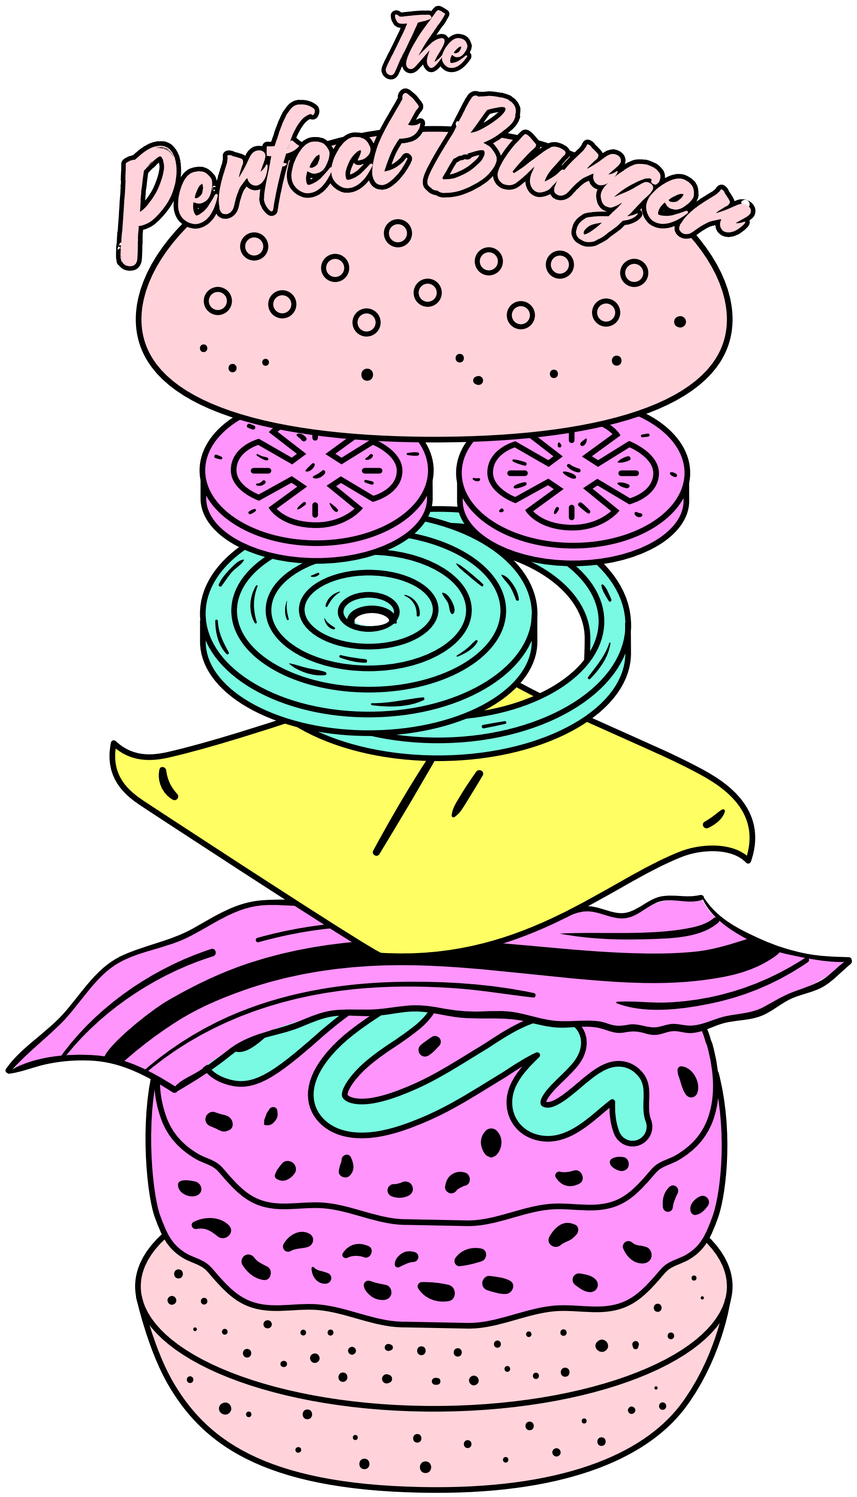

Machine Learning, the first step to AI

Machine Learning can be seen as the science of helping computers discover patterns and relationships in data. As designers, patterns form a key element in our products, so having a good learning model in place is key for success.

Basically, there are three ways in which machines can learn. You have supervised, unsupervised and reinforcement learning.

Types Of Machine Learning.

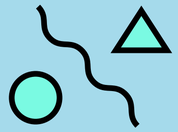

Supervised Learning.

Train an algorithm to perform classification and regression with a labelled data set.

Unsupervised Learning.

Train an algorithm to find clusters and associations in an unlabelled data set.

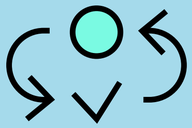

Reinforcement Learning.

Train an agent to take certain actions in an environment without a data set.

Supervised Machine Learning:

Let’s stick to the hamburger example. You want to create the best burger in the world. To do so, you first need to figure out and learn what a good burger consists of. This is where Machine Learning comes in. The machine in this case, feeds the AI with useful info like what is the best bun, best meat or best cheese to use (based on all the perfect burgers you’ve provided the machine). Machine learning in general is great in performing such single tasks.

Each part of the burger (bun, meat, cheese) will have in this case a machine learning model attached to it. To distinguish and identify what type works best per part.

And this is how it learns, as you saw you have to provide examples to the algorithm, this specific method of learning is called supervised machine learning. It works the way it implies, just like with a kid, you want it to learn but you still want to supervise the process and make sure that it is accurate. So you supervise the machine by giving examples of perfect burgers (in your view) so it can learn what it consists of.

Data that trains the model can have two forms:

1. **Labeled data** (data that is already labeled, also known as training data), labeled data is just a set of data that already has meaning, for example: *This is an apple*
2. **Unlabelled data** (data that is used to train the model)

In the training data, the model will know the features and how to properly label them and in the testing data, even if you know the label, you’ll not disclose that to the model, and that’s how to test its accuracy.

Unsupervised Machine Learning:

With this type of learning you might know which burgers are the best ones but you don't know what they are made of - the machine will cluster the dataset, that at first looks disorganized, and make its own assumptions and associations.

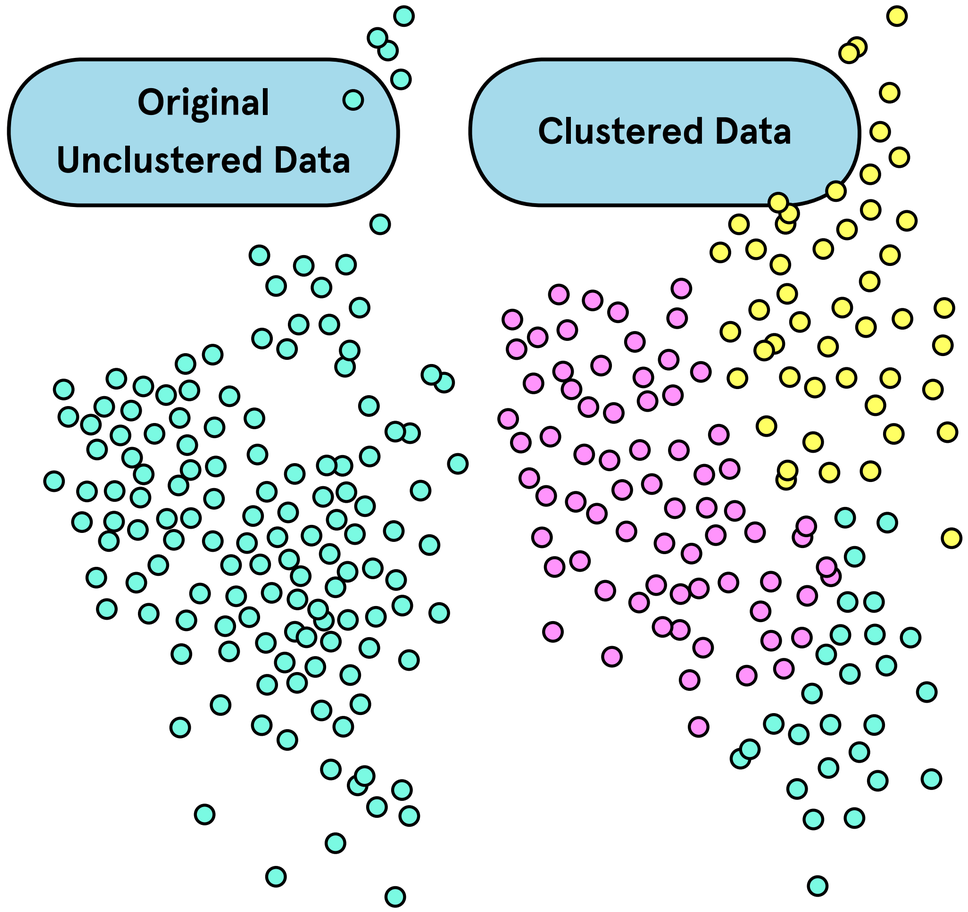

Reinforcement Learning:

Considering reinforcement learning, DeepMind is a company to follow and watch. DeepMind is pursuing the development of *General AI*. Their mission is crystal clear and they act accordingly: "DeepMind's scientific mission is to push the boundaries of AI, developing systems that can learn to solve any complex problem without needing to be taught how. To achieve this, we work from the premise that AI needs to be general".

With the help of reinforcement learning, DeepMind is developing and iterating on intelligent agents that execute reinforcement learning with unsupervised Auxiliary Tasks. Simply said, these agents perform different strategies (ways to play the game in this case) in order to find the perfect one. Everytime they get 'killed' or are 'game over', the agent performs an alternative strategy until the agent succeeds.

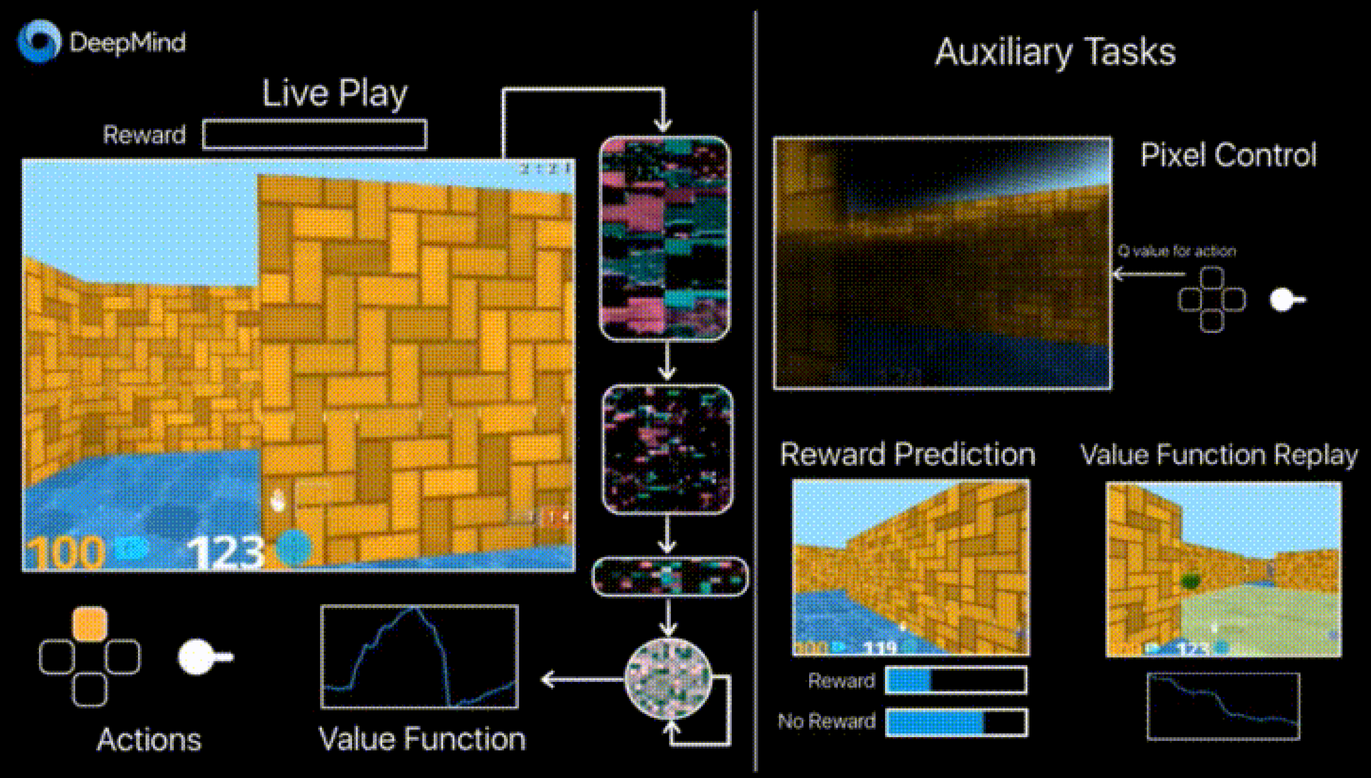

Visualisation of an agent in a Labyrinth maze.

Excited about this?

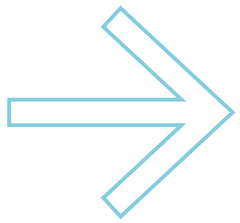

DeepMind wrote a nice paper about it.

There are several other ways to learn and all of them are based on the principle that you provide some data, you test its accuracy and you keep doing this process over time while trying to group that data in different ways in order to improve it.

Machine Learning applied on design (or as a design pattern).

Thus, now that we know HOW machines learn, let’s focus on how design adds value in this learning process.

Companies like Google and Netflix are applying Machine Learning very effectively and have an interesting approach to it. Google for example is using Machine Learning to make the mail-experience more efficient by applying automatic responses. Netflix applies Artwork Personalization to trigger user’s interest.

Gmail efficient smart reply.

Google is using Machine Learning to make the email experience more efficient and less time consuming. Google identified a need to help users to respond to an email by suggesting quick responses. Google has decades of experience in understanding languages and meanings with their Google Translate platform. They use this capability, together with a hierarchical approach to learning to predict what response is most appropriate.

Hierarchical model of Google.

Save time with Smart Reply in Gmail

Efficient Smart Reply, now for Gmail

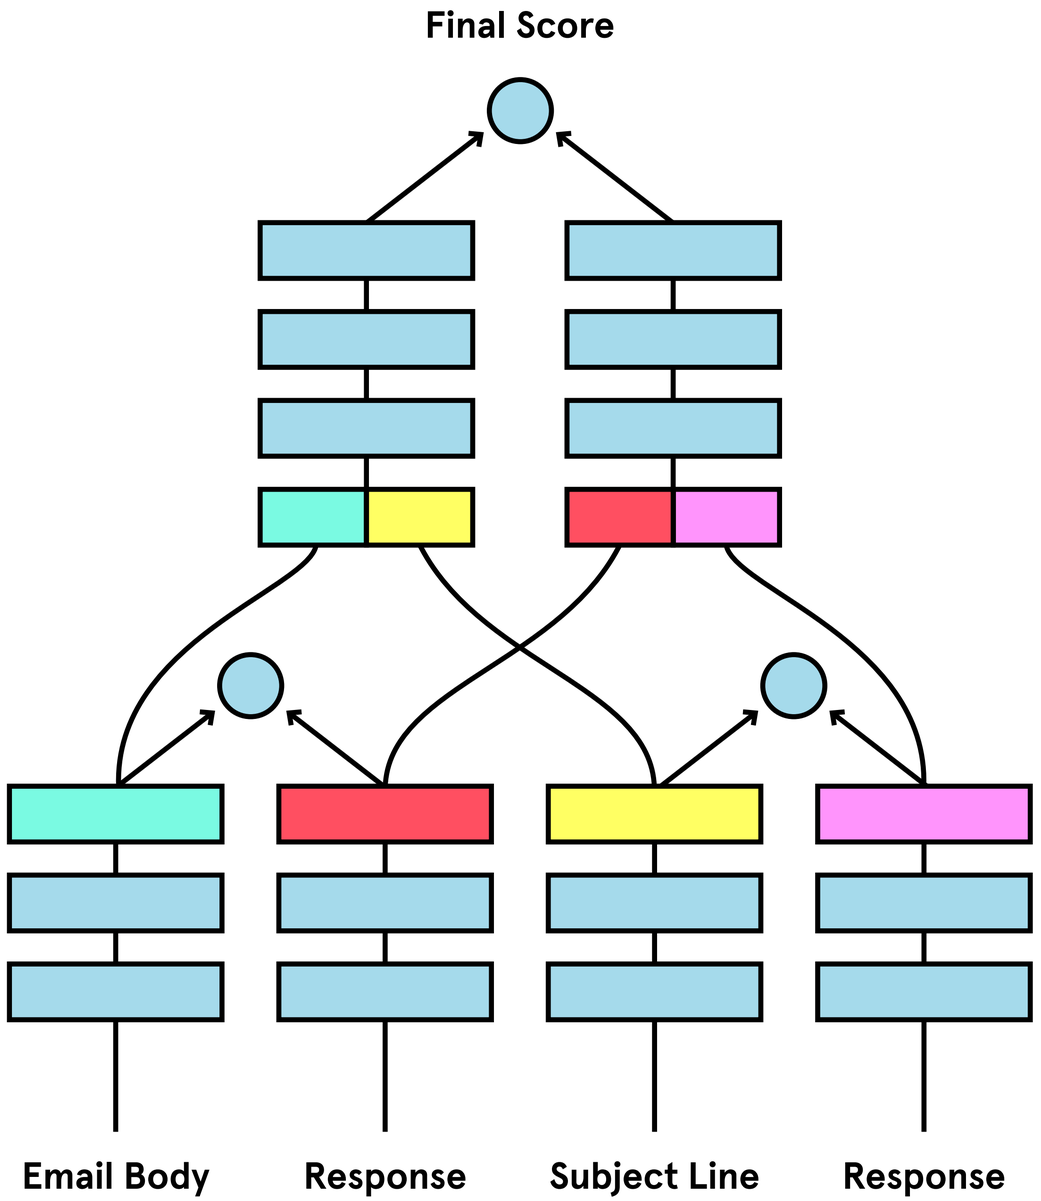

Netflix Artwork Personalization.

Most of us are watching series on Netflix and we get confronted almost daily by suggestions that are a result of machine learning – called the personalized recommendation system. Netflix wants us to discover great content. Showing only a title of a movie or series isn't enough.

It is a great effort to find the single perfect artwork for a user. How people find artwork can be very personal and based on specific taste. There is so much diversity in taste and preference that Netflix figured it would be a better idea to find the best artwork for each of their members to "highlight the aspects of a title that are specifically relevant to them". Stranger Things for example had different artwork variations, each tapping into a different theme.

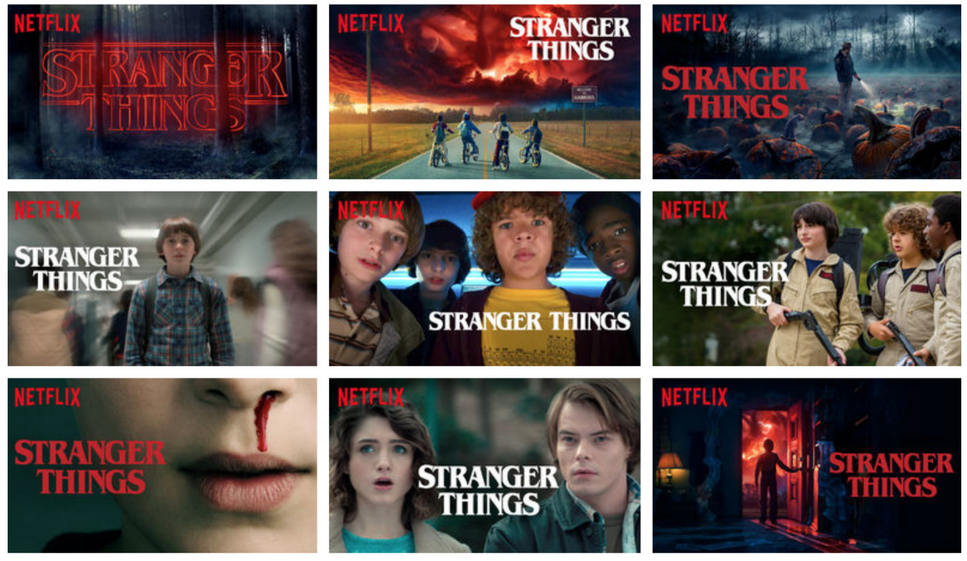

A Netflix homepage without artwork isn't very appealing.

You might wonder, how does the Netflix's Personalized Recommendation System decide what artwork to show? Your viewing history plays a big part in deciding what artwork to show. The image below shows how your viewing history affects the shown artwork based on genres you like and watch. If you watch a lot of romantic movies for example, you may be interested in Good Will Hunting if Netflix shows Matt Damon and Minnie Driver. Whereas if you're more into comedies, showing Robin Williams would appeal more to you.

The same logic can apply to actors and actresses. If you have watched many movies with John Travolta, then artwork with John Travolta would appeal more to you than one with Uma Thurman and visa versa.

Netflix's Machine Learning approach.

Much of the recommendation engine of Netflix is powered by machine learning. Netflix applies a contextual bandits approach at which they basically A/B test two algorithms against each other.

Netflix about their method:

> *"Traditionally, we collect a batch of data on how our members use the service. Then we run a new machine learning algorithm on this batch of data. Next we test this new algorithm against the current production system through an A/B test. An A/B test helps us see if the new algorithm is better than our current production system by trying it out on a random subset of members."*

This means that members in group A +get the new algorithm. If there's a higher engagement in group B the new algorithm will be rolled out — unfortunately this approach didn't present the expected results and that's why Netflix applied an online machine learning approach, to learn on the fly instead of waiting for a batch.

To learn more about how Netflix applies machine learning, check the article below!

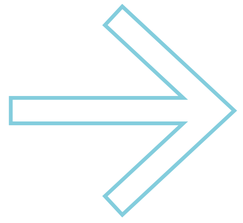

Artwork personalization at Netflix

The user on the loop.

As AI is trained with historical data in order to predict an outcome, it can become outdated really fast, take this example:

> *Let's say you build a model to anticipate which email is the most important in your user's inbox, and to do so you take a few things into consideration: Subject, Content, time it takes for the user to interact with a given email, Reply time, and other things.*
>
> *You train it extensively with historical data and now your output is 90% accurate and the model gives you a pretty good estimate, so whenever the user receives a new email, it knows how important that email is.*
>
> *Unfortunately this accuracy will not last forever because something being important is relative to innumerous things so if you don't update your model you'll be providing inaccurate predictions without even knowing.*

Adding your user's feedback into the mix might help mitigate this unfortunate outcome.

How to create this feedback loop in a ML mode?

Let's take the email example and see how Google fixed it.

> *The Gmail Priority inbox ranks email by the probability that a user will care about a given email, but something being important is highly personal so the only way to do it is by learning what is important in a per-user base and updating it as frequently as possible* ***because our priorities shift so our inbox should too.***

How is it done exactly?

Google provides new users with a *“dummy model”* and trains it every single day by actively watching how people interact with their emails. So this model will be improved every 24 hours. This illustration shows how it works.

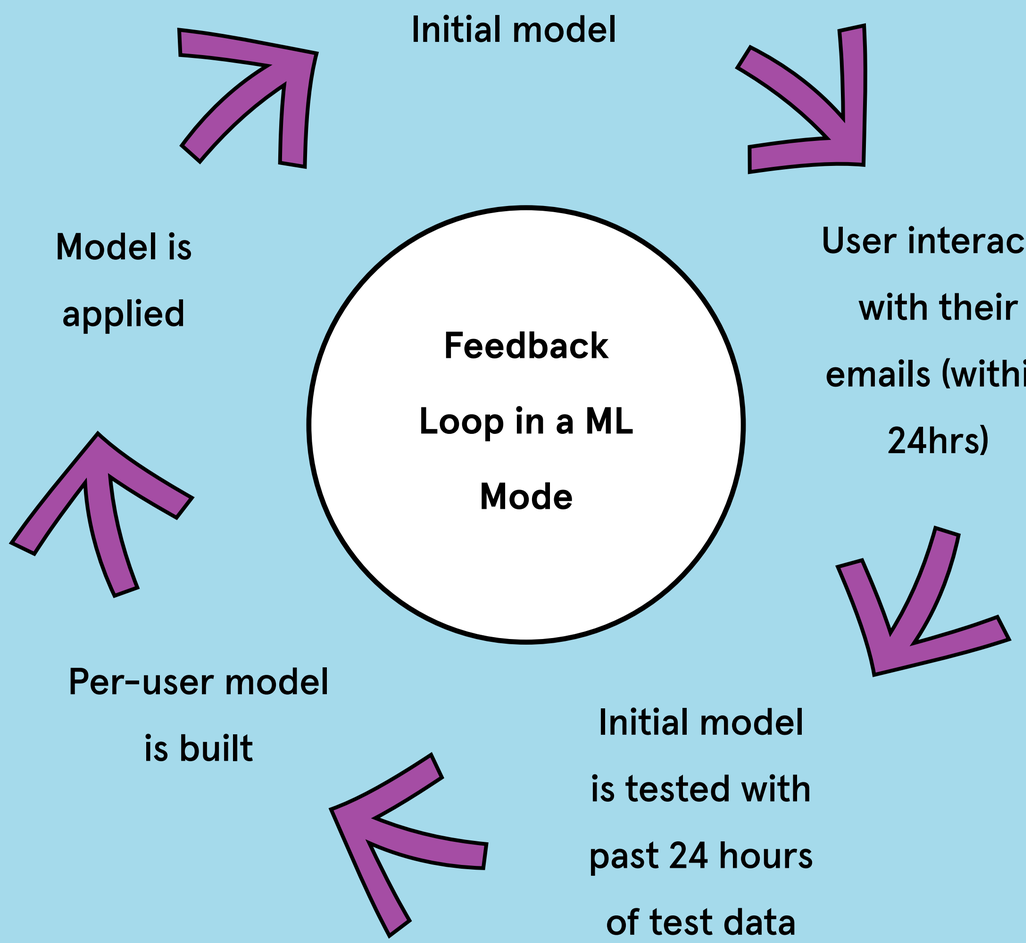

Design enabling AI.

Talking to users and taking their feedback is part of our life as designers and this is where we can contribute the most. While we're building intelligent systems we can play a massive role by making sure our user's interests and goals are what drives artificial intelligence forward.

Built-in feedback.

For example, when designing interfaces with AI we should keep in mind that it is possible that the prediction will be wrong.

A great way is to start by building interfaces that have structured feedback built in — if your model has made a mistake, structured feedback is often better than a yes/no question.

In the Netflix example mentioned earlier, they built a whole system just to try and show you the one that will make you tick, but what if you actually really dislike that artwork?

If there's no structured feedback in place, you will solely rely on what the machine says is right and you might remove its human factor from the loop.

Don't try to deceive people.

The main danger of AI, in my view, is that we might outsource important life decisions to a yes or no black box, and if we try to play the human at that moment we might make things worse.

Be sure to tell people what's happening, show them why you're recommending something YouTube does a great job at this, it recommends you new videos based on what other people have also watched:

Kevin Kenson viewers watch this

First Look at Nintendo Switch

Nintendo · 38M views · 2 years ago

"Kevin Kenson viewers watch this" with a little bit of copy you can give your users more transparency and maybe get them to trust you more in the future because now they know why you recommend them things.

Biases.

When we hear the word Bias we automatically think of something bad. It is biased.

A bias is an evolutionary gift, it is an ability to make quick decisions, a split second reaction that could save your life -- it is just a mental shortcut.

In the book *Thinking Fast and Slow* Daniel Kahneman explores the idea of a fast brain and a slow brain, the fast brain being the one that makes quick decisions and the slow one being our rational brain. The fast brain is made out biases, it is your brain quickly scanning though everything you have experienced in life and making a really fast decision.

But what do we mean when we talk about biases into AI?

There are plenty of examples of AI going wrong and being racist, sexist and so on -- how does that happen?

Let's explain it with an example:

Amazon was working on this AI to help them sort through hundreds of CVs and find the best candidate to fill a position, but as they tested the output of the model they saw that it was only selecting men and, not only arbitrarily not selecting women, but giving any female word a negative score -- Why?

When you feed a model with historical data you'll get historical bias and just like a human the computer will make a decision based on the data it has available, if that entire dataset contains mostly male words it will most likely punish CVs that contain female words.

Machines don't inherently have biases but we do and we train our models with them and there's no easy way of fixing it, most of the time we don't even know we have a bias because that's what we've been trained with.

The first step is to acknowledge that we have unconscious biases and after that, only testing and talking to your users will help you keep yourself and your product in check.

Join the ML playground!

In this chapter you've read a lot about the importance of data, how machines learn, what kind of learning methods there are and what impact biases have on AI. A lot of reading but not much interaction and fun, so that's why Pedro Marques prepared a Machine Learning playground that enables YOU to experience yourself what machine learning is and what you can do with it. Enjoy!

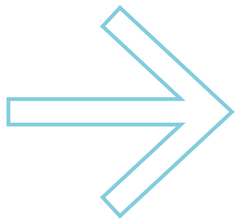

Github.

github.com/pmarquees/mlhowto

FOLLOWING CHAPTERS

This Brain food series will be released chapter-by-chapter, stretched over several months. In every chapter experts will dive deeper into specific topics related to AI.

awwwards.