

01. Are your users aware that they are interacting with an *automated solution*?

Get smarter – what is it?

Automated systems and artificial intelligence are getting increasingly better at communicating and acting like real people. They write in a grammatically correct manner and can speak so fluently that they can have conversations in ways that sound almost human. This phenomenon allows for new possibilities when setting up machines to handle communication tasks such as customer support, sales, secretarial functions, etc.

Ethical challenges arise when humans aren't aware that they are communicating with an automated system. Automated systems have become good at simulating human behavior -but they are still not

human, and therefore, they do things that are not human. They make unpredictable mistakes and are not able to understand or explain these mistakes themselves.

Humans have a need and a right to know when they are interacting with an automated system, especially if the AI makes decisions that have great importance to human life. There is a difference between an automatic e-mail confirming your online purchase and an automated refusal of the right to be with your child.



Recommendations

- Make it obvious to users that they are communicating with an automated system.
- Don't design your automated system in a way that mimics human behavior (Chatting with a robot shouldn't happen in the same interface as when chatting with a human).
- Make your automated systems more mechanical (preserve the robot voice and rigid robotic language).



The bad example

We all know Facebook's newsfeed that is used to follow the lives of our friends and family. The newsfeed has become an automated system that makes calculations on our behalf and decides which updates we should see and which we should never see. When Facebook first released this automated version of the newsfeed, few users understood this update. Many users thought the newsfeed showed them all their

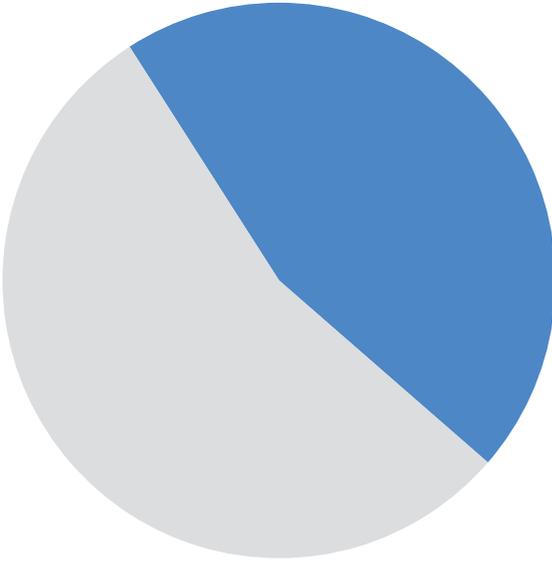
connections' updates and wondered why friends and family had nearly stopped existing on Facebook. Facebook should have made it clear that its users were interacting with an automated system. It would have given users better control and understanding of their use of Facebook.



The good example

Many media companies use news robots to write simple articles about e.g., sports scores or stock quotes. In most cases, media companies draw attention to this. Jysk Fynske Medier,

for example, writes: "Written by Jysk Fynske Medier's article robot."



02. Do your automated systems comply with *legislation and human rights*?

Get smarter – what is it?

Automated systems are increasingly being used for decision-making in cases that have an impact on people's lives.

The first ethical question you should ask yourself when designing an automated system is whether it makes decisions that respect human rights.

A major problem with automated systems is that they can discriminate unfairly in decision-making processes. Discrimination may be ethically acceptable, but not if it's based on factors such as gender, race, ethnicity, genetics, language, religion, political beliefs, disability, age, or sexual orientation (not an exhaustive list).

Also, take care to ensure that your automated system does not harm children and that it respects the right to privacy and the right to freedom of expression.



Recommendations

- Strive to ensure that your development team has the widest diversity and openness towards the minorities' use of your solution.
- Make sure that your automated system is verified by human rights experts.
- Involve minorities and vulnerable target groups in your user research.



The bad example

The Israeli artificial intelligence company Faception claims that it can analyze people's faces and predict whether they are terrorists, academics, or highly intelligent. It is, however, doubtful whether the company is capable of this at all.

There is a very high risk that their algorithms will come into conflict

with fundamental human rights regarding discrimination.

The algorithm will likely categorize people with an Arab/ Middle Eastern appearance as being terrorists.

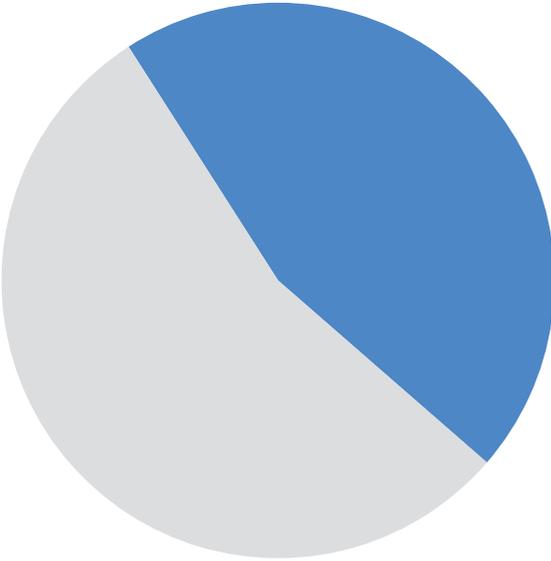


The good example

Corti is a Danish company that has developed a machine-learning algorithm that listens to emergency calls. The algorithm can recognize patterns in conversations. In this case, an ambulance will arrive more quickly.

The algorithm from Corti has learned from old emergency calls data, which means that less common dialects are underrepresented. In

the worst case, this could mean that the algorithm discriminates against these dialects and detects fewer cardiac arrests in areas where this dialect is present. Corti continuously checks the algorithms for bias, and in the case of the dialects, they have chosen to train the algorithm using data from several specific emergency calls with different dialects.



03. Does automation cause people to *lose the abilities to do a job?*

Get smarter – what is it?

You have probably tried driving a car with a GPS. You enter an address, and then the GPS tells you turn-by-turn where to go.

The GPS has meant that many people have lost the ability to read a map and have become less able to find their way without it.

Is this loss of competence problematic? It's a challenging ethical issue. But the fact is that the design of GPS' causes people to lose skills. Could one have designed the GPS differently so that people don't blindly follow the directions? Perhaps one way could be by having the north always at the top and thus helping people develop an understanding of the geography they

move around in. In other words, could one use GPS' to make people better instead of worse at finding their way around on their own?



Recommendations

- Think about the future: What would the world look like if all people lost that competence?
- Always try to design your systems in a way that doesn't make people redundant but instead makes them better and happier at doing their jobs.
- Can you incorporate learning and competence development into your digital solutions to help people develop new skills?



The bad example

Mercedes and many other car companies have developed automatic systems that can parallel park a car without the driver having to touch the steering wheel or pedals.

For many people, it is probably a great help. But it is also a feature that will mean that many people lose the ability to parallel park - an ability that can be useful for many

years to come (before cars become fully self-driving).

Should Mercedes (and others) instead have designed automated systems that help their users become experts in parallel parking? For example, by letting the drivers do it themselves, but advising and guiding underway?



The good example

Gradescope is a tool that helps teachers mark school assignments. The teacher uploads the students' work in the program, which automatically gives a grade. Then the program provides an overview of how the students are doing. Gradescope also allows the teacher to add comments and correct how students are marked, so it's not solely up to

the automation. Gradescope has thereby ensured that the automation of marking assignments does not remove competencies from teachers. On the contrary, it provides them with benefits in the form of time saved from the manual marking of assignments and an automatic overview of students' skills.



04. Is your automated system transparent, so the user can see *the engine room*?

Get smarter – what is it?

Automated systems are usually designed to make many quick decisions. These systems are generally not made for people to understand the decision-making system. In many cases, however, people must understand how an automated decision-making system works.

Algorithms can make mistakes and they can learn from “bad” data, which causes them to discriminate systematically. If you cannot open the engine bonnet to the automated system, you can’t find the errors and injustices either.

It’s only a few people that are able to open the bonnet and understand the algorithms

behind it. But transparency can also mean making one’s algorithms accessible and open to experts and legislators who act on behalf of ordinary people.

If the algorithms contain trade secrets, you can invite independent experts inside for closed ethical reviews meaning you do not have to reveal these business secrets.



Recommendations

- Try to explain to the public how your automated systems make decisions.
- Make your algorithms open and accessible so that experts can “open the bonnet.”
- Invite experts to review your algorithms.
- Explain how your automated system is trained or built.
- Try to develop simple algorithms with humanly understandable logic.
- As much as possible, avoid black box systems where you do not even understand how the algorithms work.



The bad example

In 2018, Gladsaxe Municipality in Denmark developed an algorithm to identify families where the children were likely to be mistreated. The algorithm was based on a large amount of data and machine learning that calculated the probability of children being unhappy in their families. The project met with political opposition, however. For one, it was difficult to see how the algo-

rithm decided that some children were more vulnerable than others. The lack of transparency caused citizens, politicians, and experts to lose confidence that the system was fair and sufficiently accurate. After a lot of tug of war, the project got shelved in 2019.

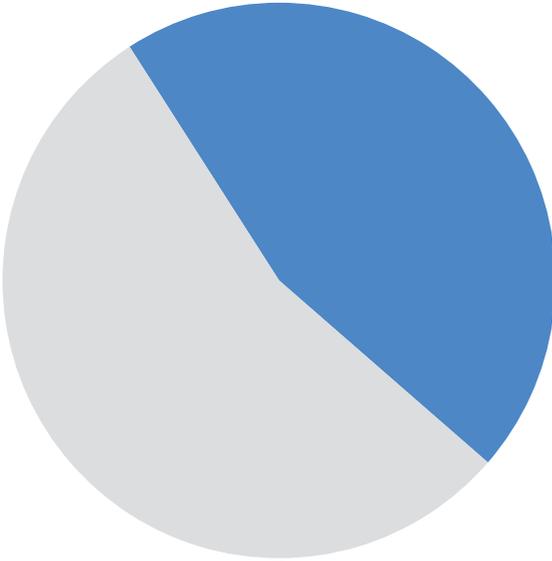


The good example

Facebook has been heavily criticized for allowing political actors to use the platform to target political ads at carefully selected segments. Using Facebook algorithms, actors can hit select groups with perfectly designed messages that the public will never witness. To address this issue, however, Facebook has made political ads open to the public. Anyone can view all the ads from a particular political actor, see how

much they spend on advertising, and roughly who they are targeting with their ads.

The system is not aimed at ordinary people but rather at journalists who can use the system to monitor political parties and actors, typically during elections. The system is not perfect, but it's an excellent example of implementing openness into an automated system.



05. Can your automated system *explain itself*?

Get smarter – what is it?

As citizens of democratic and free societies, we get explanations for the decisions that affect our lives. “You must pay a reminder fee BECAUSE you paid three days later than the agreed deadline” or “You must serve 30 days in prison BECAUSE you violated section 266 of the Penal Code”.

However, many automated systems are designed in such a way that they are not able to provide such explanations. And especially machine learning systems can be bad at providing justifications that humans can understand. It can create ethical issues in cases where people expect to get an explanation for an algorithmic decision, but the system is unable to provide one.

As automated systems make more and more decisions, there is a need for the systems to be designed in such a way that they can explain their decisions. However, it is largely an ethical trade-off when a system should be able to explain itself. Of course, this applies to decisions with far-reaching consequences for human life, but in practice, there will be many algorithmic decisions that do not need explanations because they are too trivial and mundane (for example, an algorithm that automatically turns off the light in a room).



Recommendations

- Always try to design algorithms that explain themselves while they're being used.
- If possible, avoid black-box algorithms where an explanation is important.
- Always make sure through the digital design that users can request an explanation an algorithmic decision.
- Ultimately, one should always be able to get an explanation from a human being for an algorithmic decision.
- Make your algorithms open and accessible so that experts can "open the bonnet."
- Invite experts to review your algorithms.



The bad example

Facebook has employees whose jobs are to view content that is reported by users or by algorithms. But due to the amount of content on the platform, the algorithms can take complete control in certain situations, for example, during the covid-19 pandemic, when most of Facebook's employees have been home. It has turned out that the algorithms have been the cause of

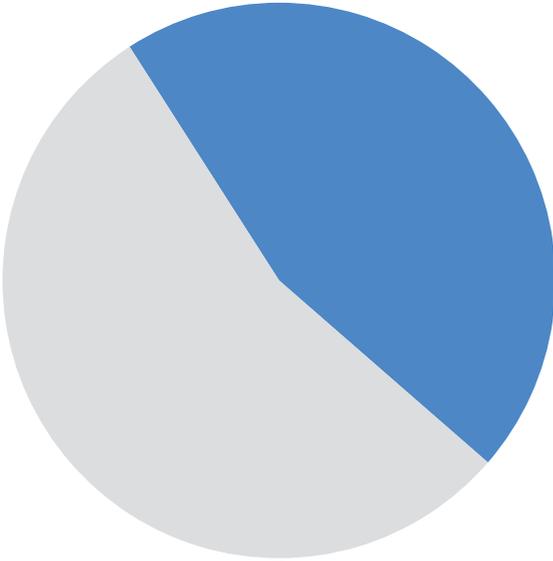
much content on the platform has been deleted, or profiles have been reported without proper reasoning. Craig Kelley (an MP in the UK) has found that his posts have been deleted without explanation. After complaints, Facebook could not explain why they had been deleted and denied responsibility, but it turns out that this happens for about 300,000 Facebook posts a day.



The good example

Rainbird is a company that provides algorithmic decision-making systems to financial companies. Rainbird helps banks and insurance companies respond to customer inquiries and detect fraud. Rainbird differs from other algorithmic systems because it always incorporates explanatory systems so that

customer service employees can understand how the algorithms arrive at a decision. It means that they can give their customers human and understandable explanations if their account has been closed or if they have been denied a loan.



06. Are your algorithms *prejudiced*?

Get smarter – what is it

Modern artificial intelligence and automated systems use machine learning and data from our society and world.

The ethical problems arise firstly when these data are bad and do not represent the real world, e.g. if one only trains a face recognition algorithm on white people. Secondly, biased algorithms can arise if they reflect a reality that is already biased. The algorithms simply automate existing biases such as racism or gender discrimination. One talks about that the algorithms become biased - that is, skewed so that they do not represent reality or so that they discriminate in a way that is not desirable.

It is important to understand that one can never remove biases completely, so the goal is not bias-free algorithms but rather algorithms where one knows their biases and where their biases are in line with widely accepted human biases.



Recommendations

- Verify that machine learning data represents all the stakeholders that may be affected by the algorithm.
- Avoid having your digital solution automate existing and unwanted biases.
- Regularly check your algorithms for bias and preferably use independent experts for this.
- Make sure to have diversity in your development team to raise awareness of unwanted bias.



The bad example

IBM, Microsoft, and Mevii have developed facial recognition software that they claim can identify people with 99% accuracy. But a study by MIT Media Lab found that this accuracy was only valid for facial recognition of white men. The accuracy decreased when identifying women and black people and

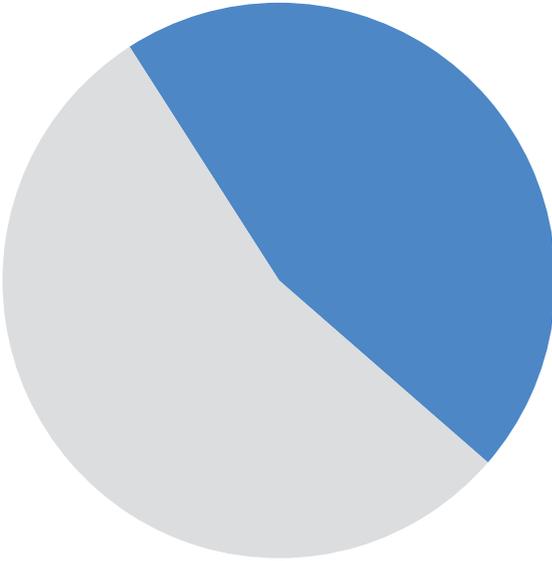
mostly with black women, where the accuracy was 65%. It turned out that the datasets for this software were images of parliamentarians, which may explain the difference in precision.



The good example

When companies write job advertisements, they can, through their use of language, discriminate against people based on gender, age, and social background. One example is writing in a language aimed at young people rather than older people. Textio is a program

that uses artificial intelligence to identify this type of discrimination and guides companies towards more inclusive language use. Companies can thus use Textio when recruiting to ensure more diversity among employees.



07. Is there an *unnecessarily high risk* with your automated system?

Get smarter – what is it?

Automated decision-making can be divided into four categories with different ethical risks:

1. The system makes precise decisions on issues with little to no consequences. Here there are no considerable risks.
2. The system often makes erroneous decisions but on issues with little consequence. There is a case to be aware of risks here, but the effect of failure is minuscule.
3. The system makes precise decisions, but the consequences of errors can be fatal. Here one should take great care to be sure that the automated system is robust and

not biased. One should consider whether there should be used automation in the first place.

4. The system often makes erroneous decisions on issues with fatal consequences. Here one should always avoid automation!



Recommendations

- Consider the consequences if your automated system fails.
- Work with worst-case scenarios.
- Be sure to monitor and evaluate the errors that your automated system makes.
- Involve independent experts and critics in the development of your automated solutions.
- Be careful with untransparent black box systems used to automate critical functions.



The bad example

IBM's artificial intelligence, Watson, is used to assist the examination, diagnosis, and treatment of patients. In 2018, it was discovered that Watson recommended incorrect and sometimes deadly medica-

tions to patients. This discovery resulted in Rigshospitalet in Denmark and Novo Nordisk abandoning their use of the technology.

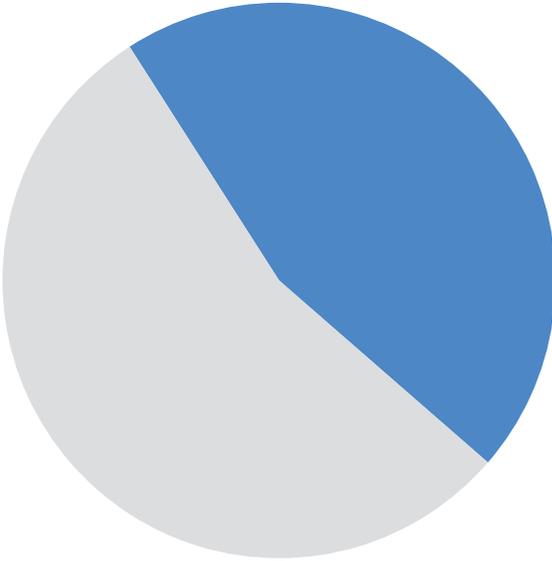


The good example

When an app for digital coronavirus contact tracing in Denmark was in development in the spring of 2020, a broad societal and ethical debate arose about automated monitoring of people's contact with each other. There was an important argument from the critics saying that the value of the contact tracing app did not measure up with the risk of the government gaining access to citizens' location data.

The risk of abuse was too big.

The solution was to involve an expert group with an understanding of both ethics and technology, which resulted in a final solution where the risk of data misuse had been minimized by decentralizing data collection. The project could have been completely abandoned, but instead, it was decided to design a solution that minimized risks.



08. Is someone in the company ready to step in *when automation fails*?

Get smarter – what is it?

Artificial intelligence can seem superhuman and infallible because it can find patterns and make calculations on data otherwise incomprehensible to humans. But artificial intelligence also makes mistakes, and often these are surprisingly banal because artificial intelligence lacks the human understanding of how the world is connected.

It is therefore crucial that humans are never removed from automated decision-making systems. Firstly, you should make sure that you have someone constantly keeping an eye out for errors and irregularities in the system. Secondly, you should always ensure that your customers and users can get in touch with a human being if your

automated system fails. The latter is also a requirement of the GDPR.



Recommendations

- Always be very careful if you remove people completely from your automated systems.
- Include a functionality in your digital solution that allows users to speak with a person.
- Make sure you have the right people on standby to intervene when automation fails.
- Be aware that working as a supervisor of automated systems can be monotonous, tedious, and uncomfortable.



The bad example

In 2011, the American t-shirt company, Solid Gold Bomb, developed an algorithm that could make funny slogans about the meme “Keep calm and carry on” for print on t-shirts. The system was fully automated, so the funny ones were put up for sale on Amazon without any people having checked the slogans. Solid Gold Bomb didn’t deem it necessary to observe the t-shirt slogans because the t-shirts would only be produced if people

bought them.

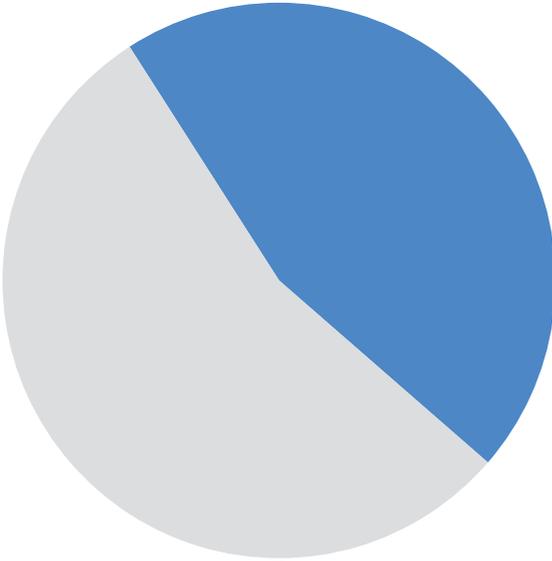
Unfortunately, the algorithm began putting t-shirts up for sale with messages like “Keep calm and kill her” and “Keep calm and rape a lot.” The case exploded on social media, and the company went bankrupt due to bad publicity.



The good example

The Danish company Holo works with self-driving buses, which are currently driving around Copenhagen, Aalborg, and Oslo. The buses can only get up to a speed of 15 km/h and have never been in any

accidents, but Holo has still chosen to place a living person in all their buses, ready to intervene if the bus fails or if any other problems occur in the bus.



09. Is your automated system *adaptable to changes?*

Get smarter – what is it?

Automated systems are usually designed based on historical data. One automates actions and workflows that work today and reckons that they will also work in the future, but the world is constantly changing.

Humans are changing preferences, attitudes, and patterns of action, which means that most automated systems that interact with humans will stop working if they are not continuously updated.

There can be several ethical consequences of static automatic systems. A self-driving car that does not get its algorithms updated with new maps will drive the wrong way. But static algorithms can also perpetuate

unwanted prejudices that may have been acceptable in the past but have changed over time, for example, discrimination against women in the workplace.



Recommendations

- Always consider your automated system as a work-in-progress. It is never a finished product.
- Always train your machine learning systems on new data.
- Be aware that your automated system may need whole new types of data.
- Be sure to design curious artificial intelligence that tests and searches for changes and new patterns.



The bad example

In 2009, based on millions of users' searches, Google Flu Trends managed to track down a flu epidemic in the United States two weeks faster than the U.S. Center for Disease Control and Prevention. Due to this, the expectation became that Google's algorithm could predict pre-

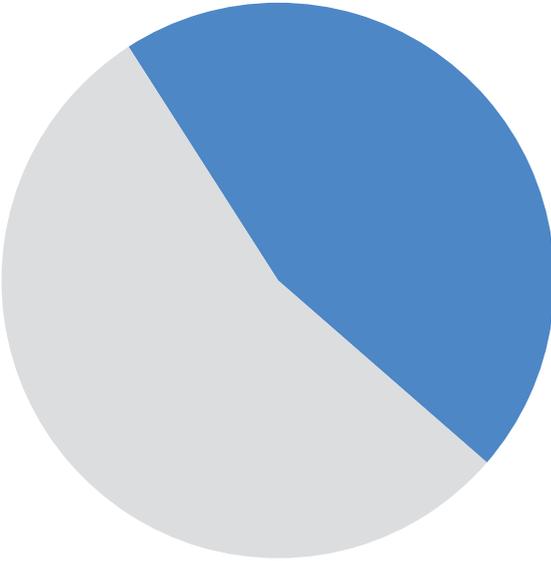
cisely where and when an epidemic would strike in the future. But after 2009, Google's predictions became inaccurate, and on several occasions, the service overestimated the scale of the outbreaks to such an extent that Google Flu Trends went extinct after five years.



The good example

When listening to music on Spotify, the company collects data about users' taste in music used to generate music recommendation lists. It could develop into an echo chamber where the user listens to some music, gets recommended more of the same, and keeps listening to the same kind. But Spotify's algorithms are good at testing the limits of their users' musical taste, which means

that users are constantly allowed to listen to new music and test their limits. The result is that the algorithm both helps push people's music taste and follows when people change it over time (unless they are stuck on old eighties songs).



10. Can your automated system *be hacked?*

Get smarter – what is it?

Self-driving cars use artificial intelligence to interpret the sensory impressions they encounter through the car's many sensors.

The problem is that these sensory impressions can be hacked without breaking into the car but simply by changing the surrounding environment. For example, people have found that you can make self-driving cars overlook stop signs if you stick white tape in specific patterns on the sign. It is self-evident that you need to secure your digital solutions against traditional hacking, where people break into a system. But modern machine learning, which is based on real-world data, allows for entirely new ways of hacking.

Often it is not even malicious hacking but simply people wanting to make fun of the "robots." For example, when people jump out in front of self-driving cars to test whether they will brake.



Recommendations

- Think in worst-case scenarios: There will always be someone trying to cheat your automated system.
- Consider how people will react to your automated system and take into account their reactions.
- Be aware that unethical systems will create more motivation for hacking and data manipulation.



The bad example

A Vietnamese IT security company has demonstrated that they can hack the facial recognition feature on an iPhone X. This is done by making a 3D printed "twin mask" for less than 2,000 Danish kroner

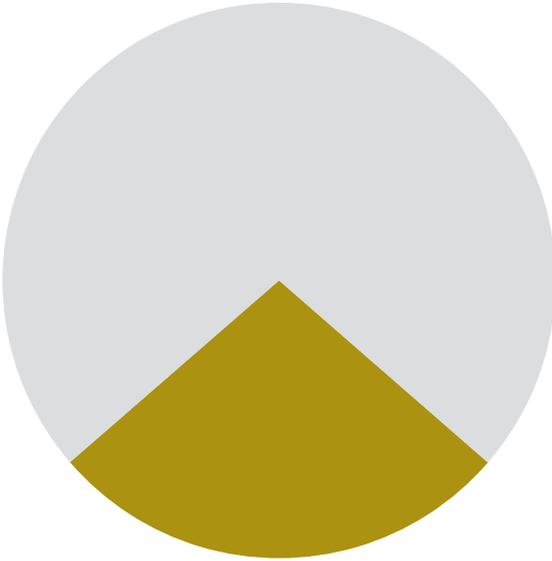
(around 300 euros). So you do not have to be an expert in computers to break into an iPhone -all you need is a picture of your victim and access to a 3D printer.



The good example

Google is the world's most important search engine, and it can be life or death for businesses whether they are at the top of Google's search results or not. Therefore, of course, many try to figure out Google's algorithms so that they can get higher up on the results page. Sometimes people cross the line and use methods that try to hack Google's algorithm. It is also known as black hat search engine optimization.

Google is in a perpetual battle against these hackers, who do not break into Google's systems but try to hack the data that Google uses to rank results. Google has so far won the battle against the black hats. But it is an eternal battle that requires thousands of dedicated employees who are constantly developing their algorithms.



01. Does your design play with *negative emotions*?

Get smarter – what is it?

Studies show that most people will go out of their way to avoid a loss compared to what they would do to achieve the same gain. If you want people to do something, then it is proven more effective to scare them with something negative than to motivate them with positivity, a tactic used by sellers, insurance companies, and baby equipment retailers for years.

In the digital world, motivation through fear becomes even more effective because, combined with user data, it can make fear messages increasingly personalized.

There is, of course, an ethical balance. It is fair game when a pension company informs

people that they are not saving enough for their retirement, but if they combine this message with an image of an old lady eating cat food, you start moving beyond the ethical limits of digital behavioral design.



Recommendations

- Think about whether you use fear, uncertainty, and doubt as motivation in your design.
- Be especially careful with fear messages based on knowledge about the user.
- Think about whether you can turn negative communication positive.
- Be aware that false information such as “only one product left” may be false.



The bad example

Many hotel booking websites use messages that tell the user that x number of people have already booked and that you need to hurry to book at the price shown. Users are scared that the offer and dream trip will disappear or that the price will double any minute. This practice is in itself

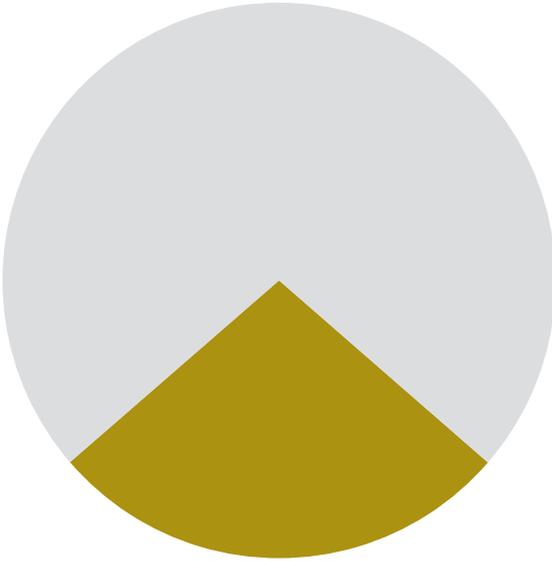
unethical because it uses fear and insecurity to motivate users. But it can also be outright illegal if the stated deadlines and bookings are not truthful. Many booking sites have been in the spotlight of the competition authorities because of this practice.



The good example

Tobi is a Danish startup that helps parents invest their children's savings. Their message is that it is far more efficient to invest their child's savings than to leave it in a regular bank account. They use comparisons to show what this difference can mean when the children turn 18, but they do not use fear in these messages.

They use concrete examples, but they do not tell stories about how horrible it will be to be 18 years old without child savings in a world where a two-bedroom apartment in Copenhagen will cost 10 million Danish kroner (or well over a million pounds or euros).



02. Do you deliberately make it difficult for users to find or understand *information or functionality*?

Get smarter – what is it?

Digital design is most often about creating interfaces that are user-friendly and useful for users. But companies can have an interest in making users do things that are not to their benefit. For example, making users continue to subscribe to a service that they do not use, buying more products than necessary, or perhaps saying yes to conditions that are not in their best interest.

In these cases, design can develop into a manipulative design. Also known as a dark pattern where you use cunning design tricks to get users to do things that are not to their advantage. Often it is about hiding important information or highlighting other information using colors, animations, font

sizes, or other graphic tools.

The ethical guideline must be that you should always design based on your users' interests. There is, of course, a balance to be found. But one should always avoid deliberate manipulation, where one can say with certainty that it is about benefiting the company at the expense of the user.



Recommendations

- When designing, you should not only think about usability but also about user needs.
- Consider how you would feel about using your own digital design.
- Pay special attention to children and other vulnerable audiences who are particularly easy to manipulate.
- User test all the things that are not for the benefit of the company (unsubscribing, returning products, etc.).



The bad example

Do you have an Amazon account?
Try deleting the account!

If it seems tricky, you can find a manual here: kortlink.dk/2agfy

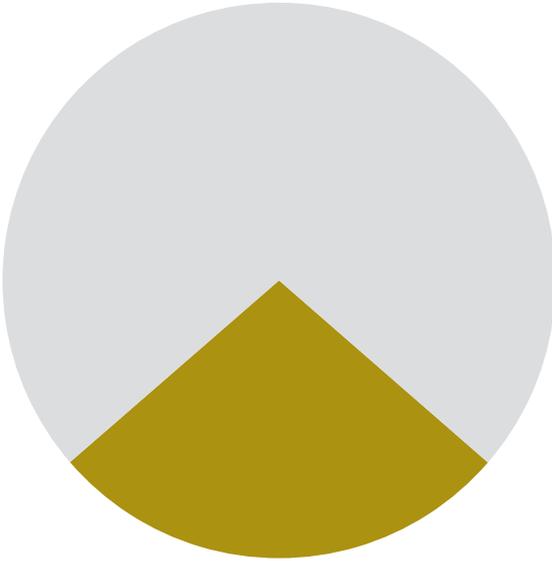
Hint: You need more than ten clicks and it is actually not possible without contacting someone from Amazon.



The good example

The vast majority of subscription-based digital services offer a free trial period with the condition that you provide your credit card information so that the subscription can start automatically at the end of the trial period. The result, of course, is that many users forget to unsubscribe and therefore remain stuck as paying customers, even though they may not have wanted to. However, there are exceptions - namely the Danish newspaper Information and the cycling training platform Zwift. Both of these offer a free trial period where you do not have to provide your card information.

Both stand out from the competitors and appear more ethical and more attractive to customers.



03. Do you exploit your user's *inability to concentrate* to your own advantage?

Get smarter – what is it?

Users can easily read long texts and familiarize themselves with complex issues online. But it requires that they are in the right mindset, such as when they are listening to a podcast or reading long magazine articles.

However, if the user is in the process of a transaction or buying products, downloading apps, or the like, then one cannot expect users to have the capacity to familiarize themselves with complex issues about cookies, tracking, use of data, profiling, etc.

In other words, you should not expect users to have understood your data policy or your

subscription rules just because they clicked the accept button. You may have received a legal acceptance, but ethically, you have not received an actual endorsement.



Recommendations

- Avoid long and complex texts as much as possible.
- Try to split information into smaller chunks and present it when relevant.
- If you need acceptance from the users, ask for it in the context where it is suitable.
- Consider whether you could make your service less complex.
- Use skilled copywriters to write this kind of text.



The bad example

The University of Copenhagen has a privacy policy that is certainly not transparent to the general public. For example, it states that the University of Copenhagen:

"(...) registers and processes personal data based on Article 6 of the

GDPR. The processing of sensitive personal data in research projects is covered by (...)" One must have a legal understanding and be familiar with the various sections of the GDPR to understand how the University of Copenhagen processes data.



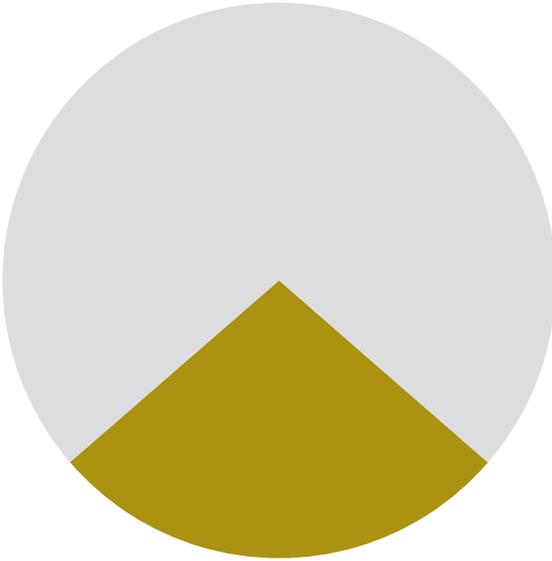
The good example

The DuckDuckGo search engine declares that the platform is not storing, sharing, or using personal data about the users.

They sum up their privacy policy with the phrase: "We don't collect or share personal information. That's our privacy policy in a nutshell." Next, it is explained in human language what data is collected and how. It is still a lengthy document but

purposely made to be as understandable as possible.

The example also shows that data-ethical companies typically find it easier to formulate comprehensible privacy documents simply because they have less to hide and more to be proud of concerning their customers.



04. Do you *manipulate actions* by taking advantage of *people's need to be social*?

Get smarter – what is it?

People are social, and most of us are constantly looking for social recognition from our surroundings.

It is an urge to be exploited for better or worse. On the one hand, our need for social recognition means that we help each other with advice and guidance through social media. But this very urge can also be why teens spend thousands of dollars on digital skins that make them more popular on Fortnite.

In the digital world, companies can exploit our urge to be accepted socially to such an extreme where the consequences can be digital dependence or over-consumption of

digital products or services.

Ethically, it is okay to use social design techniques to motivate people to take certain actions, but it is crucial to keep your users' deeper interests in mind. Think about whether your social interfaces make people happier or more unhappy.



Recommendations

- Beware of using social actions as a currency that users need to invest in to achieve something else. Social actions should be the goal in itself.
- Consider how social designs can develop when a lot of users use your service. Does the service change character when everyone uses it?
- Embed stop-blocks in your social designs so that people do not get carried away. For example, set restrictions on purchases or time consumption. Keep in mind that social interaction can be very addictive.



The bad example

When Pokemon Go launched in 2016, it was a pretty innocent augmented reality game that was all about collecting digital Pokemon monsters in the real world. The game assumed a social character as people began to share information about the locations of the best Pokemons, which meant that big cities suddenly were overrun by

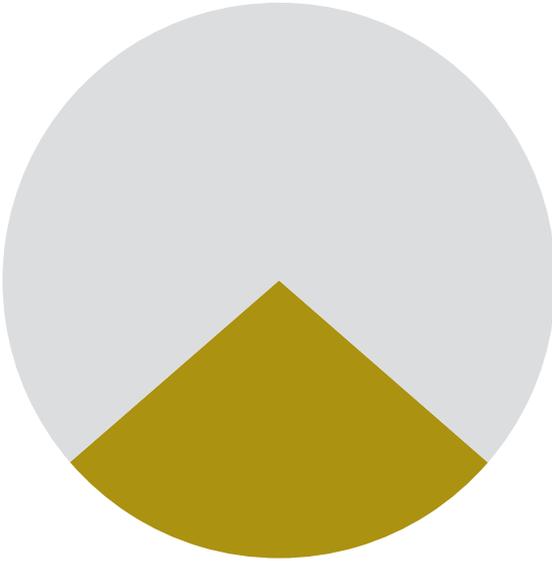
swarms of thoughtless Pokemon hunters who looked down at their screens and trampled everything and everyone down on their way. It was not a planned consequence but a consequence of the game becoming socially addictive and lacking stop blocks that could prevent large numbers of people from gathering.



The good example

Instagram is, if anything, a service that benefits from the urge for social recognition. But in 2020, after much criticism, Instagram admitted that it could go out of hand for some users, who became heavily dependent on reaping as many likes as possible. To reduce the reliance on social recognition, Instagram removed the ability to see how many likes photos have received.

A relatively small initiative on a large platform like Instagram but a positive example of a small step in a more ethical design direction.



05. Are you trying to *create an addiction* to your product with cheap tricks?

Get smarter – what is it?

Many companies want their products used as much as possible, especially if the products are financed through advertisements.

It is therefore tempting to use nudging, social mechanisms, or data to get their users hooked. It can mean excessive use of notifications, or it could be the use of features where people increase their social status by being more active on a platform. Of course, it is ethically okay to design products to be so good that people become addicted. It becomes unethical when the addictive tricks solely contribute to creating addiction and do not make the product itself better. It can be a tricky ethical balance to strike!



Recommendations

- Ask your users if they feel they are spending too much time or too much money on your service.
- Include stop blocks that prevent excessive dependence on your product.
- Think of vulnerable target groups such as children or gambling addicts who become more easily addicted.



The bad example

Snapchat has a feature called streaks, which motivates people to keep Snapchat conversations running for as long as possible. The longer the conversation runs back and forth, the better a streak and the more rewards the conversation partners receive in the form of funny

emojis. A classic example of a design that is not making the product better but solely serves the purpose of keeping people on the platform.

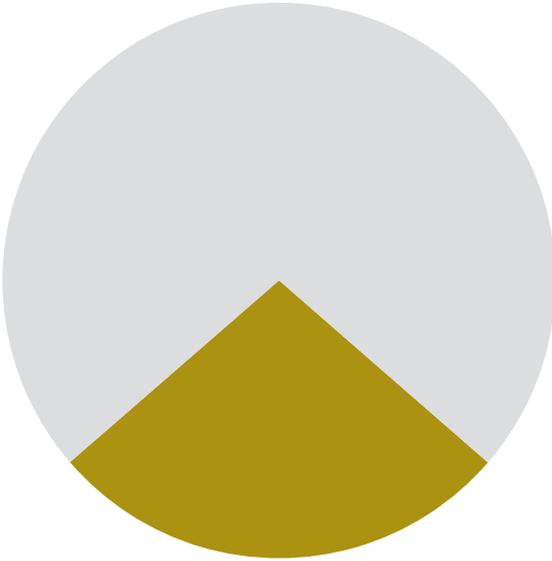


The good example

Netflix is a service that can be very addictive because you always get recommendations for relevant content based on your personal preferences. However, Netflix also had a feature that automatically initiated the next episode of a series once having finished. Here the users were being nudged to get stuck and binge many episodes in a row.

It is an example of a design that does not give the user much extra value (it is not difficult to click on the start button), but in return, it can create an addiction where the user does not benefit.

Netflix has since removed this design element, so the series no longer start automatically.



06. Do you *validate* or *challenge* your users?

Get smarter – what is it?

The customer is always right! Companies must always give customers what they want, and with artificial intelligence and algorithms, it becomes easier to figure out precisely what customers want.

The problem is that you can end up knowing your customers so well that you never challenge their preferences but give them more of the same.

Consider whether good customer service is about more than giving customers more of the same. Perhaps it also entails challenging customers and showing them surprising new products or information.

It is a difficult balance where one must not become patronizing and manipulative. People want to choose freely, but most also have a strong desire to discover new opportunities and see new perspectives on the world.



Recommendations

- Build algorithms that do not create echo chambers.
- Build algorithms that deliberately make random mistakes to expose users to less of the same.
- Consider your metrics: Make it a KPI to expose users to new content or new products.
- Combine - as far as possible - algorithmic recommendations with human recommendations.



The bad example

If you observe young children using YouTube, you will notice how they are quick to click on the videos recommended to them by the algorithm. A child can quickly go from watching an informative video about dinosaurs into a meaningless universe of cheaply produced cartoons intended for children to click on. The algorithms lure children into

echo chambers of content that have the sole purpose of making the children linger (and watch commercials).

In other words, YouTube is an unsuitable medium for children, which is bad for both children and YouTube.



The good example

Most music streaming services make use of algorithmic recommendations, which create personalized playlists for their users. Many of these tend to create musical echo chambers where the users' musical tastes are rarely challenged. In this market, the Tidal service stands out

because they have hired a small army of curators who make human-made playlists. Not all users like Tidal's playlists, but the strategy gives Tidal an ethical profile that stands out significantly from the competition.



01. Are you collecting *too many* data points, and do you keep them for *too long*?

Get smarter – what is it?

When data is considered the new oil and artificial intelligence become better with more data, it can be tempting to collect data without restraints.

helps to strengthen data inequality between businesses and citizens.

However, you should be aware that excessive data collection is illegal under the GDPR, and besides, it is also unethical, impractical, and risky. Large amounts of data increase the risk of data leakage, and it also makes it more difficult for you to handle data on behalf of the user.

Finally, it is profoundly unethical to collect data that is not highly business-critical because it exposes users to unnecessary risks of privacy breaches and because it



Recommendations

- Always try to remove a data point for users rather than adding a new one.
- Always ask yourself if you need this particular data point.
- Be sure to clean up old data that is no longer needed.



The bad example

Danish high schools use the platform Lectio for administration and communication between teachers and students. The problem with Lectio is that the data collected is not deleted, which means teachers can go back in time and find information about grades and sick leave on students who have long

since finished high school. The data is available so that all teachers at a school can access the information without a password.

The practice is probably a violation of the GDPR, but it is also unethical handling of data because they are storing data that serves no purpose.



The good example

DuckDuckGo is an alternative to Google's search engine that does not collect user data or track user searches on the web. DuckDuckGo does not even know how many users they have, as they do not track users.

Because DuckDuckGo does not store information about users, users will only see ads related to their current

search. When using DuckDuckGo, it is clear to the users that they aren't being tracked across web pages and that the search engine is not collecting too much data about their movements online.



02. Do you *anonymize* your data?

Get smarter – what is it?

By default, you should always assume your data collection could be published online tomorrow. What would the consequences then be? Can individuals be identified in your data? Or have you ensured that the data is so anonymized that nothing can be disclosed about the individual?

Anonymizing is not easy. It is not enough to remove names, addresses, or social security numbers because someone's identity can often be deduced by combining or merging data with other publicly available data sets.

You may not collect personal data at all, but you should be aware that seemingly harm-

less data can also help identify individuals. For example, one can derive dietary habits from users' recipe collections to determine any chronic diseases.



Recommendations

- Ensure that data about people and their identifiable information are securely separated.
- Consider whether you really need data about individual people or if aggregating data is sufficient.
- Invite independent experts to assess whether your anonymous data can be de-anonymized by merging data.



The bad example

In 2017, the training app Strava published a world map, which showed aggregated data for all the cycling and running routes that people had uploaded to Strava - seemingly hassle-free, because all data was anonymized.

But the problem was that in countries like Afghanistan, there was very

little activity apart from selected locations, namely US bases, where soldiers used the app when training. It turned out that the map came to show precise locations and routes of bases and soldiers who exercised on their own in the landscape. In other words, it was showing sensitive data!



The good example

Using the Strava training app to track and share bike rides seems innocent to most users. But few people think about the fact that knowledge of where a bike ride starts and ends is also knowledge of where an expensive bike is parked. Precious data for bike thieves! Therefore, Strava has chosen to obscure the exact point

where a bike ride starts and ends, even though it provides a worsened user experience for both those who share the bike ride and for those who follow their friends' bike rides. They have sacrificed a little on the user experience in exchange for giving users a much more secure app.



03. How do you store data?

Get smarter – what is it?

There is an ethical and legal duty to store data securely.

Nevertheless, errors still occur when storing data, and the errors are often due to companies and organizations simply not thinking about storing sensitive data.

Digital platforms often make it very easy to collect and store data. Therefore, data collection can also happen without control within an organization. An example is the Danish company Medicals Nordic, which was responsible for testing corona patients. The company used WhatsApp for daily workflows, and when they suddenly had to test thousands of Danes, they chose to

continue using WhatsApp as the platform where they shared test results. Subsequently, the company was fired by the Danish regions. But the example shows that it is often simple thoughtlessness that leads to insufficient data processing.



Recommendations

- Make sure that all employees are familiar with the fundamental rules regarding the use of data. (Facebook, WhatsApp, Dropbox, etc.).
- Consider your use of online platforms where you do not have complete control over data.
- If you are not an expert in data storage, you should hire external experts to secure your solutions.



The bad example

Copenhagen Zoo has 140,000 annual cardholders. The log-in page for these annual cardholders had no restrictions on log-in attempts, which made it easy for unauthorized people to try out and access the cardholders' personal information, including card number, name, address, and e-mail. It got reported as a breach of personal data security and was criticized by the Danish Data Protection Agency.

Now, the zoo has had their members change passwords and they have introduced the function "I am not a robot," which ensures that a program cannot cheat the system. Also, they have introduced a feature that, after three failed attempts, the system will block access for one hour.



The good example

Telegram is an online chat service similar to WhatsApp and Facebook Messenger. However, Telegram has a strong focus on encryption and privacy protection. Telegram has, among other things, a data-safe function called Secret Chats.

When sending messages via Secret Chats, the chat is fully encrypted, and no data is stored centrally

or is accessible by employees at Telegram. You also cannot forward Secret Chats, and you can even set messages to self-delete for both the sender and recipient after a certain period.

Secret Chats are therefore safe, as long as you have your phone safely stored in your pocket.



04. Do you give people *access to their own data*?

Get smarter – what is it?

If people are to have control over their data, they also need access to their data. This means:

- That you must ensure that all data you have collected about your users is visible to the user.
- That even though you collect data from many sources, you should make sure that the user can access their data from one place.
- You have a responsibility to present data so that it's understandable to the user, even if it may be cryptic.

You are ethical if you have a dedicated area in your digital solutions where people can access data that pertains to them. It must be easy to find, and the data must be easy to see and understand. If data is used for creating new data (for example, through profiling), then the new data must be just as clear and easily accessible. Also, people need to have control over data. Essentially, users should always be able to delete data. If it makes sense, users should also have the option to correct their data to make it more accurate.



Recommendations

- Always consider data as your user's property
- Help your users/customers to understand their data rights
- Make sure to communicate about data in a way, so the user understand it is their property
- Ask your users how they would like to access their data.
- Test your solution to see if users can access and understand their data.



The bad example

Most Danes use an online banking solution where you usually have full access to all data about your finances. However, banks often use financial data to profile their customers in different earnings segments, and very few banks exhibit this profiling to their customers.

Banks also have data on how much they earn on each customer (the price of being a customer in the

bank), but this information is also not readily available on online banking platforms.

Banks are thus adept at providing raw data, but when it comes to aggregated data, which can be very valuable to their customers, their digital solutions are severely limited.



The good example

Following the Cambridge Analytica scandal, Facebook was criticized for collecting too much data about its users and doing so without users' awareness. Subsequently, Facebook has designed a comprehensive and user-friendly area on the platform where you can get an overview of your data. Facebook also provides the options to restrict the collection of data and delete

personal data. Because Facebook's business model is about data collection, the company often makes it difficult to find places where you can restrict their access to data. In some cases, users get warned that they will lose functionality if they don't allow the data collection, which is both unethical and on the edge of EU law.



05. Have you obtained user permission to *collect and process data*?

Get smarter – what is it?

Most companies have understood that it requires permission to collect data about their customers and users.

The ethical problem is more often about whether one has obtained real consent where the user also understands what it entails.

Think of the many cookie pop-ups that users encounter online today. How many users know what they are doing when accepting cookies?

You have an ethical obligation to obtain permission in a way that is understandable to your users. It often means that you have

to communicate much more concisely and pedagogically.

But remember that if it is difficult to explain why you are collecting data, then there may be a case for not collecting it in the first place.



Recommendations

- Do not let lawyers write the texts for the users/customers alone. Let communicators write it.
- Remember that rarely are texts of more than 5 - 10 lines read all the way through.
- Consider having two documents: a formal/legal one and a document that is easy to read but perhaps not entirely legally correct.
- If you cannot get actual and fully informed permission from your users, do not collect data.



The bad example

Most Facebook users are aware that their data is used for targeting advertisements and other content. But only the rarest of users on Facebook understand how their data can be used and misused.

A glaring example is the Russian-developed app Girls Around Me. The app combines freely available data from Facebook and

Foursquare to create a stalker app where (typically) men can log in to the app and view women nearby with data about when they were last located at a specific location.

The app caused quite a stir when it launched and has since been removed from the Appstore.



The good example

At the time of writing, Danish telcom uses a cookie pop-up that gives users a choice between "All cookies" and "Only necessary cookies." With a single click, the user can deselect all cookies except the technical cookies necessary to make the site work.

It is a user-friendly way to get acceptance for cookies. You even have the opportunity to dive deeper into the

information and adjust your choice further (which, however, is probably only done by a few).

On the other hand, TDC has chosen to color the "All cookies" button an alluring blue, while the "Only required" button is white on a white background. Here, there is no doubt about what TDC wants from the user, and the company could have designed it more ethically.



06. Do you inform your users about *how they are profiled*?

Get smarter – what is it?

Artificial intelligence and algorithms are getting better and better at finding patterns in data, and often they see patterns in data that are not immediately visible to the human eye. Companies and organizations can develop profiles about their customers and users that contain knowledge that not even the users themselves know. Companies can determine people's creditworthiness and preferences in books alongside mental illnesses, sexual preferences, and political attitudes.

It creates some obvious ethical challenges, both concerning inequality and human rights. Profiling can establish a high degree of inequality between the company that

profiles and the person who (perhaps unknowingly) is a victim of profiling.

Profiling can also conflict with fundamental rights, which are about not storing sensitive personal data or discriminating based on gender, race, sexuality, etc.



Recommendations

- Be aware that database profiling can violate fundamental human rights.
- Do not create profiles that end up being personal data.
- Always inform your users and customers about how they are profiled.
- Make it easy for users to understand how you have created their profiles.
- Profile data is also the property of users, and they have the right to access and delete this type of data.



The bad example

Target is an American retailer that collects large amounts of data on people's buying behavior, using this data to profile their customers. The profiles are for sending tailored offers on products. One way is by creating profiles on whether customers show signs of being pregnant to send suggestions on pregnancy and baby products. In 2012, this

resulted in Target sending pregnancy-related offers to a young high school girl, even though neither she nor her parents were aware of her pregnancy. This story created problems in the small family and made Target reconsider its use of profiling in marketing.



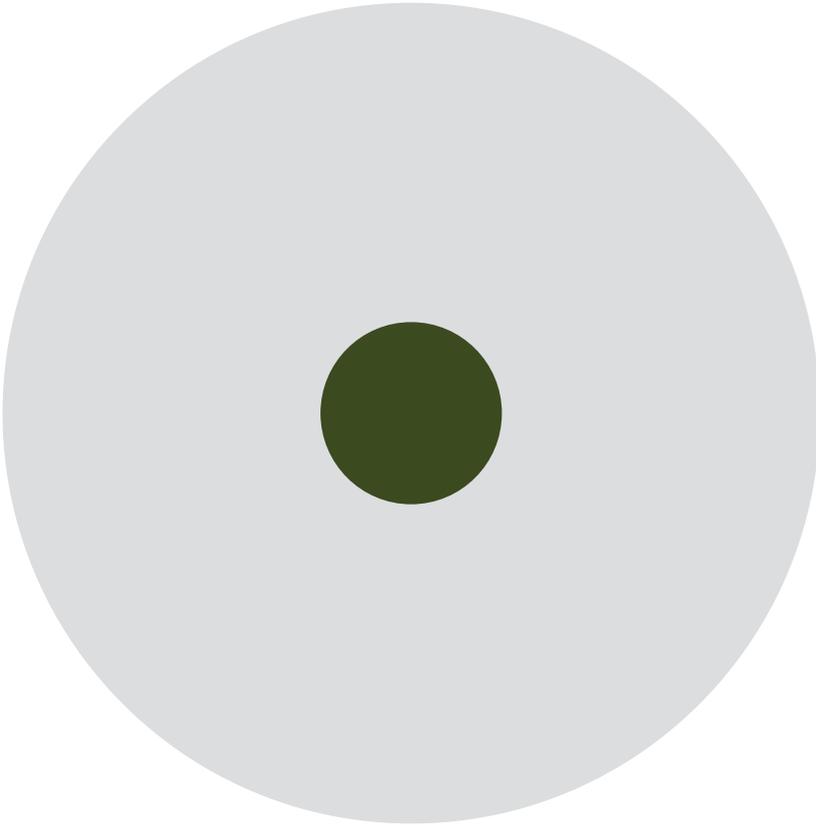
The good example

Facebook is rarely a company showcased as particularly data ethical but, the company is good at showing how their collected data categorizes (profiles) the individual user in different areas of interest, used to target ads.

hundreds of interests linked to your profile on Facebook.

Users can also choose to delete information to remove personal marketing within this subject area.

You can find this information under privacy settings on your Facebook. It's easy to find and user-friendly. You can get an overview of the



Put *the human* in the center

Get smarter – what is it?

If you only need to remember one thing about ethical digital design, it is to put humans at the center. It sounds simplistic, but often the human being, whom it's all for, can slip into the background when designing digital solutions. Remember the human, and remember that humans are not just abstract persons, but living beings of flesh and blood - just like yourself. And remember that the human in this situation is not only your users or customers but all those who may be affected by your digital solution.



An example

You are in the process of developing an app to help people exercise more. You have developed personas and tested your app on the core target group. You have applied all the principles of user-driven design.

But have you remembered to involve people with disabilities who also want to exercise but who have difficulty with many traditional

exercises? And have you thought of the future consequences? Can situations arise where people overconsume your product and get sick from exercising too much? And have you ensured that users have control over their data? In other words: Did you remember to put the human in the center?



Avoid manipulating

Get smarter – what is it?

Often, digital solutions are designed to help people by making decisions on their behalf. But beware that this help does not undermine people's very basic urge to decide over their own lives and take extra care to ensure that your digital solution does not end up manipulating people in a harmful way.



An example

You are selling cycling clothes online, and now you have designed a well-thought-out Facebook campaign targeting people who are interested in cycling.

If they have visited your website, you use retargeting to hit them continuously with targeted messages on Facebook. You can see from your statistics that your new

targeted Facebook campaign is 77% more effective than a regular campaign.

All your competitors do the same thing, but you ask yourself if it's okay to use data, behavioral design, and algorithms to manipulate so many people's buying behavior.



Make your technology *understandable*

Get smarter – what is it?

Digital solutions are often complex and difficult for ordinary people to understand. But that's why you still have a responsibility to make your digital solutions understandable and transparent to let users comprehend how the solutions work and affect their lives.



An example

Your lawyers have asked you to post a 30-page “terms and conditions” document on your website. All users must approve the terms and conditions to be allowed to use your website. But you can see that the legal text is challenging for your users to understand, so you work with the lawyers to cut the 30 pages down to five pages in understandable language. But you soon find out that your users do not read these five pages either. What will you do?

You can avoid many legal issues by collecting and exchanging less data (which you did not need anyway). You, therefore, end up simplifying your solution, which leads to better conversion rates. Additionally, you should strive to explain the consequences of users’ choices in the interaction design itself, instead of hiding the explanations away in a cryptic document that users never read anyway.



Avoid creating *inequality*

Get smarter – what is it?

Digital design without thought for ethics can often end up perpetuating and reinforcing existing inequalities in society. Always think about designing solutions that do not create more imbalances.



An example

You are in the process of designing a new digital insurance product, where the insurance company can track the quality of their customers' driving behavior. If customers drive well and carefully, the premium will automatically be lower.

80% of all Danes think that they drive a car better than average, so your product will naturally become popular very quickly.

But have you considered the consequences for those who are bad at driving and suddenly get very high premiums? And are you sure that your algorithm is good enough to distinguish bad drivers from good ones? And have you considered the societal consequences of more and more insurance products becoming so specific that the collective element of insurance is starting to disappear?



Give users *control*

Get smarter – what is it?

Digital solutions may well help people and make their lives easier, but they must not leave people with a sense of losing control. Always make sure to design solutions that give people more and not less control.



An example

You are in the process of developing a GPS for cycling, where exercise cyclists can plan their trips and ride their bicycles along different routes. You are designing a user-friendly interface where cyclists are only told when to turn right or left. It's easy, and even people who are bad at reading maps can figure it out. But user tests show that people quickly turn off their brains and lose track of the landscape they are in. If the GPS

runs out of battery, they are lost and have no idea where they are.

Therefore, you change your solution a little by showing a map that always faces north. It's a bit more complex, but you can quickly see that cyclists are happier with the overall experience because they feel they are in control of their bike ride.