



**College
of
Health
and
Medical
Techniques**

Data Science Ethics



**Al-Mustaqbal
University**

**Stage 2 , Semester 1
@ Department of Intelligent Medical Systems**

**Prepared by
Dr. Samir Badrawi**

History, Concept of Informed Consent

The majority of this course material is based on Coursera

<https://www.coursera.org/learn/data-science-ethics>

“H.V. Jagadish lectures”, a Professor at the University of Michigan

1. Human Subjects Research and Informed Consent:

Part 2

In traditional experiments on human subjects, you have what is called **prospective data collection**. That is, you set up the experiment and you collect the data from this experiment that you have set up.

1. review by the IRB, to get approval.

(this approval is required whenever there are human subjects, whether it's medical research or social science research).

2. do the experiment and collect data.

Cont.

this kind of process is enforced by US federal law for any government funded experiments

The process **prospective data collection** is not just applied to *medical experiments* and *social science experiments*, but also to *computer science experiments*.

So, if one is conducting a **user study** of a new **user interface**, for example, for a web program, that if it is a *research study* . . . it requires IRB review.

However, there is an exception in this which is a very important exception in terms of the law :-

The IRB review is required only if it is **research**, and it is not required if it is **ordinary conduct of business**. What does this mean?


If you're a **web company**, you probably do something called **A/B testing**. This kind of testing is used not just by web companies, but it is almost always used by companies that are on the web because they have high customer volumes.

Cont.

The kind of thing that *these companies* do is take some fraction of the site visitors and show them version **A**, and they take the remainder and show them version **B**.

A/B Testing

- Routinely used
- Particularly by companies with high customer volume
- Particularly for web-page design



Cont.

And then they'd watch the reactions of people shown version **A** and people shown version **B**. And if it turns out that people shown version **A** behave in a certain way, let's say click on a link that generates revenue **29%** of the time versus people shown version **B** only clicked **13%** of the time, then version **A** wins this **A/B** test.

And from that point on, they would only deploy version **A** of their web page design.

Cont.

Companies in the *Information Technology industry* conduct **A/B** tests all the time.

In fact, large companies like *Google* are said to be conducting hundreds of **A/B** tests simultaneously at any point in time.

Clearly *human user are impacted*, and clearly this is an *experiment*. The company doesn't know whether the **A group** or the **B group** is going to do better.

Cont.

However, this is *not considered human subjects research* because the company is actually trying to do the best thing in terms of whatever it is that it's trying to maximize in choosing **A** and **B**.

Companies are showing **A** and **B** as choices because they think that **A** and **B** are both reasonable choices and they'd like to see which one people actually prefer.

Cont.

Similarly, companies understand that humans are not always rational and that there are psychological cues that can push our buttons. And they definitely try to push our buttons to buy more or to click more.

And this is considered *good business practice*.

Human Subjects Research?

- **Web Companies do A/B testing all the time. Human users are impacted.**
- **Companies try to “push our buttons” to buy more.**
 - *This is considered good business.*
- **Companies can have imperfect algorithms that give us poor results.**

Cont.

Companies also are imperfect. Algorithms are used for many things and these algorithms are in most cases *imperfect*. They give poor results sometimes.

But again, *we as customers of these companies, as subjects to these algorithms*, just live with these poor results.

So all of these things are just best effort companies doing what they always do, and there is not a problem with these being human subjects research.

Cont.

In contrast, consider the following incident that took place in in 2012 :

Researchers at *Facebook* and *Cornell University* conducted an experiment where the news feed of selected Facebook users was manipulated:

some users were shown more *positive articles*, others were shown more *negative or disappointing or sad articles*.

Cont.

What they (*FB & Corn univ.*) found was people shown more positive articles posted more positive articles themselves on *Facebook*. People shown more negative articles posted more negative articles themselves. In other words, they demonstrated ‘*emotional contagion*’.

And this was something that was considered an important result psychologically and a paper describing this experiment and the results was published in the Proceedings of the *National Academies of Science* in 2014.

The users of *Facebook* **did not know** that their newsfeed was being manipulated.

They **did know** that *Facebook* regularly tinkers with the algorithm to choose items that are shown in their newsfeed.

And they **did know** that what they actually saw in their news feed was whatever *Facebook* thought was good for them to see.

What they **didn't know** was that what *Facebook* claimed was good for them to see was potentially manipulated for the purposes of this research study.

Cont.

Once the results of this experiment became public, *Facebook* got a lot of bad press.

It came out of this experiment is that companies should not lie to us intentionally, even if what they're doing in terms of the effects on us are not that much different from if they were just conducting regular business.

Cont.

There's a company called OkCupid, which is a matchmaking company that also conducted experiments reporting compatibility scores that were different than they had actually estimated in their algorithm.

And they actually thought that they were doing the right thing.

So, Christian Rudder, the CEO of OkCupid, wrote a blog post saying, "*We experiment on human beings in support of what Facebook had done.*".

Cont.

What we learn from these examples is that when something new is possible such as high volume experiments on human subjects across the web, that we have to feel our way through to a social consensus on what is appropriate to do.

Clearly, somebody like Christian Rudder thought that they were doing something that was appropriate at the time, at which they did not.

Cont.


Today, there is good consensus that companies should not conduct this kind of research even if they're legally allowed to do it.

It's one of these things where we have an ethical consensus that it is not right for companies to lie to us in conducting research.

3. **Limitations of informed consent**


The concept of **informed consent** was developed in the context of experiments that would be conducted on human subjects, and data would be collected prospectively after a consent had been obtained.

Let's look at how this translates into our world today and the kinds of things that data scientists do.




First, the experiments are not often experiments, there is first **collection of data**, and the experiment comes afterwards.

And the **collection of data** is often from people who are interacting with **somebody** who would like to collect that data, possibly a merchant, possibly a **software vendor** of some sort.




Informed, in this case, is usually something that's hidden in multiple pages of fine print. You want to use some service, and you're given a multi-page record full of dense legalese. And you're required to say, I accept, before you can actually use that service.

There is some benefit in the law for the person who's getting this kind of consent, but it is very far from being something that is clear, firm permission.




And from an ethical basis, setting aside the law, we all can say that there is some weirdness if we claim that somebody has been informed, because they were given multiple pages of fine print that they really didn't have an opportunity to read.

The **notion of voluntary** also is a little questionable. Because the **consent** is being obtained exactly at the time that the user is intending to perform a particular action, like use a software service or buy a product




This is not something that they have time and opportunity to think over, this is not something that was shown to them early during their shopping experience where they could worry about what was required by different vendors, and they could fold this into their choice of vendor, or choice of product or service.

Rather, this is quite late in their decision-making process, this is after they've already decided what they really want to do, and now suddenly it is, pay me this toll or you can't go on this road, after you've decided to drive a particular route




If one looks at how **informed consent** works, and think about the example of what Facebook does, and how one would think about a research experiment that Facebook may conduct to do, say, some psychology study.

Facebook explicitly tells the user in its agreement that it may collect user data for research purposes. Certainly, it has been doing so since 2012 after its famous (**mood contagion**) experiment.



The issue here is not that there's something wrong with Facebook. Facebook actually does a very good job of thinking issues of privacy and what users would like to do with their data.

It's just that Facebook is at the head of the curve. It's a big company and it has a lot of very personal data for a lot of us, and often gets into the cross-hairs of the community before anybody else, because they're there before anybody else.



Okay, so that was informed consent with regard to data collection. But then there's a question of, what is the data actually going to use for?

So, I may give data about myself to a merchant to obtain a specific service.

I don't want the merchant to use these data for other purposes.

I don't want them to use it to sell me other things, for instance.

I may want them only to use it for the specific service that I have contracted with them for.

Repurposing of Data !

Repurposing Can be Problematic


- I give data about myself to a merchant to obtain a specific service.
- I do not want the merchant :
 - *To use these data for other purposes*
 - *To share these data with others*
- **Consent can often be limited to disallow repurposing.**
- **“Context” matters**



And I don't want the merchant to share these data with other users.

So I may, for instance, give them a consent to disallow repurposing. And this is where we're saying, you may collect this data, but the data that you collect is something where you've been given permission in a particular context.

That context matters and you may not use it in any other context or sell it to somebody else.




Now **repurposing** of data is not all bad, there are often business needs to do this. So my credit card company obviously needs to collect data about my purchases and about my payments.

They don't have to share this data with the credit reporting agency, and maybe I don't particularly like that they share it with the credit reporting agency.

Repurposing is not all bad

- **I understand that my credit card company collects data about my purchases and payments, and I am willing to accept that they also share some of this with a credit reporting agency, even if I don't particularly like it.**
- **(Ideally) the card company tells me about this intended sharing, and I (may be reluctantly) agree.**



But this is something that I've got to accept as a part of the social setup, and this is something that the credit card company ought to be telling me about and that I agree to.

If I'm going to be given credit, I need to participate in the ecosystem that involves credit reporting. And a separate credit reporting agency learning things about my purchases and payments from my credit card company.




Repurposing, even if it isn't of business necessity, might actually be of great societal value.

I share medical data with my hospital to get better medical care. But I may not actually mind at all, I might be very glad if my medical data is repurposed for medical research.

I might feel happy that information about my disease progression, my health, is going to help future generations beat some scourge.


Repurposing is often necessary

- I share medical data with my hospital to get better care, but I may gladly support its repurposed use for medical research.
- However, the specific research questions may not be known at the time I receive my care.
 - *Retrospective Data Analysis*
- Can we define a consent agreement broad enough to cover a range of possible medical research without being so broad as to cover “everything under the sun”?



The thing is, the specific research questions that would be asked, the things that scientists wish to study, may not be known at the time I receive my care. The questions come afterwards.

The data was already collected, and this is called **retrospective data analysis** as opposed to **prospective data collection**.



And the problem here is, how can we obtain a consent agreement where we've given enough information to the human subject that they know what they're consenting to? and yet have the consent be broad enough to cover a range of possible research questions that one might wish to ask.

And so, this is a balancing act in terms of having an informed consent with enough information for it to be meaningful.

To conclude, most data of interest, most data that we'll analyze are created by humans, are about humans, or have impacts that affect humans.

Humans and Data

- **Most data of interest are:**
 - *By humans*
 - *About humans, OR*
 - *Affect humans*
- **So we have to consider this impact as we practice Data Science**

When we practice data science, we have to consider this impact, and it is consideration of this impact that is at the cornerstone of ethical practice of data science.

4. Case Study: It's Not OKCupid

Matchmaking companies exist on the Web. OkCupid is one of them.

The way these things work is customers sign up, enter a profile, and the company tries to find compatibility between user profiles based on what you've told about yourself.

You'll be recommended people to meet that the company thinks might be compatible with you.



Have an account?

Sign in



Join the best free dating site on Earth.

I am a

Gay

Woman

Continue

☐ I don't want to see or be seen by straight people.



Signing up takes two minutes and is totally free.



Our matching algorithm helps you find the right people.



iOS or Android?
You can take us to go.

"The Google of online dating" - The Boston Globe

About OKCupid | Legal | © Humor Rainbow, Inc.



31% Match
16% Friend
82% Enemy

DanTheMan412

26 / M / Gay / Single
Ann Arbor, Michigan

Message

Remove Favorite

Hide Report

About

The Two of Us

Personality

Tests

My self-summary

Hi, my name is Danny and I am an english teacher. Although I am well read, I really enjoy abstract poems like this one by John Rives:


O}-<Q<=10018MTWThF
02 28 08 7.55pm 1* *10/-

His Details

Last Online December 27


Ethnicity White

Height 5'9" (1.75m).



Determining compatibility, is of course, extremely difficult, and a company like OkCupid has a proprietary algorithm to try to predict this compatibility.

They do so by means of developing a compatibility score that is computed through some proprietary function of all the different things that you have in your profile.



Of course, they're trying all the time to improve their algorithm and improve the predictive power of the compatibility score that they create.

And for this purpose, they're constantly monitoring how well they do. They're constantly tinkering with the algorithm, trying to improve it, change some weights, change some parameters.

And some of these changes may be successful and may be adopted. Others may be rolled back.

And so, in this sense, if you're a user of a website like OkCupid, you're constantly being experimented upon in that what you're experiencing is an algorithm that is perpetually in development.

Now, OkCupid actually took this thing a couple of steps further. They tried a slightly bigger experiment



So, for example, a thing that they started to do was something called **the Love Is Blind Day**.

And on this day, they would *suppress photographs* in user profiles. And what this meant was that you would not be able to see the profile photographs of potential people to meet.

You would of course be able to read what they had written in their profile in terms of their interests and other text attributes.



31% Match
16% Friend
82% Enemy

DanTheMan412

26 / M / Gay / Single
Ann Arbor, Michigan

Message

Remove Favorite

Hide Report

About

The Two of Us

Personality

Tests

My self-summary

Hi, my name is Danny and I am an english teacher. Although I am well read, I really enjoy abstract poems like this one by John Rives:

O}-<Q<=10018MTWThF
02.28.08 7:55pm I*_*IQ<=
O}-<1/2?...1/256?oma!♥

His Details

Last Online	December 27
Ethnicity	White
Height	5'9" (1.75m).

What they found was that even though, as humans, we all very much care about photographs, and that is likely to be a large part of the input in terms of our deciding if somebody looks interesting and attractive, that not having those profile pictures actually resulted in more conversations, more dates.

And in that sense, their platform was actually more successful on this day, on ...

the Love Is Blind Day.



So, taking this kind of counter-intuitive success another step further,

OkCupid decided to try an experiment where they wanted to understand the impact on the success of a date of simply being told that you are compatible.

And so what they did was took people who were actually not very compatible, had a low compatibility score, and **falsely told them that they had a high compatibility score** just to bring people together.

Conversely, they also took people who had a high compatibility score and **told them they had a low compatibility score**.

And they actually did a three-by-three matrix of this sort with 30, 60, and 90% compatibility score, actual versus 30, 60, 90% declared.

And what they found was the success rate was the **highest** when people **were compatible** and were **told they were compatible**, that is at the 90-90 entry of this matrix.

And it was **lowest** where they were **not compatible** and were told they were **not compatible** at 30-30 entry of this matrix.

That if people were **not compatible**, but were told they **were compatible**, *they actually did pretty well.*


And in fact, they did just as well as if they were compatible but were told they were not.

This is an interesting insight into human behavior. And this experiment was reported by the CEO of OkCupid in a blog post.

When he reported this, he was skewered on the internet by a whole range of people.


And fundamentally, what this says is in doing this experiment interesting, as it might be, that *OkCupid broke the unwritten expectations of society*.

The CEO, Christian Rudder, obviously didn't agree with this general-social consensus or maybe there wasn't enough of a social consensus that he and other people like him would think it perfectly acceptable to do what they did in terms of reporting compatibility scores that were intentionally incorrect.



Really what we're doing here is socially drawing a line between experiments that companies conduct in the normal course of business because they're trying to improve their product or improve their algorithm.

On the other hand, having a company intentionally lie to you, intentionally give you a wrong score, is something that most people would consider socially unacceptable.



And because of this particular experiment and the discussions that followed, we say that this is a matter on which there is greater social consensus today than there was at the time that OkCupid actually did conduct this experiment.