


Social media advertising for clinical studies: Ethical and data protection implications of online targeting

Big Data & Society
January–June: 1–15
© The Author(s) 2023
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/20539517231156127
journals.sagepub.com/home/bds


Rainer Mühlhoff¹  and Theresa Willem² 

Abstract

Social media advertising has revolutionised the advertising world by providing data-driven targeting methods. One area where social media advertising is just gaining a foothold is in the recruitment of clinical study participants. Here, as everywhere, social media advertising promises more yield per money spent because the technology can better reach highly specialised groups. In this article, we point out severe societal risks posed by advertising for clinical studies on social media. We show that social media advertising for clinical studies in many cases violates the privacy of individual users (R1), creates collective privacy risks by helping platform companies train predictive models of medical information that can be applied to all their users (R2), exploits the weaknesses of existing guidelines in (biomedical) research ethics (R3) and is detrimental to the quality of (biomedical) research (R4). We argue that the well-intentioned promises, which are often associated with the use of social media advertising for clinical studies, are untenable from a balanced point of view. Consequently, we call for updates of research ethics guidelines and better regulation of Big Data and inferential analytics. We conclude that social media advertising – especially with vulnerable patient populations – is not suitable as a recruitment tool for clinical studies as long as the processing of (even anonymised) social media usage data and the training of predictive models by data analytics and artificial intelligence companies is not sufficiently regulated.

Keywords

Targeted advertising, clinical studies, predictive analytics, data ethics, data protection and privacy, biomedical research ethics

Introduction

With the rise of social media, new communication possibilities have emerged, amongst them also targeted advertising. All kinds of offers are being communicated through *Facebook*, *Instagram*, *TikTok*, *Twitter*, *Grindr*, *Craigslist* and other social media platforms, making social media one of the most used advertising channels. At the same time, the failure to recruit an appropriate cohort for a clinical trial is an ongoing threat to the efficiency and effectiveness of clinical research (Gelinus et al., 2017; Puffer and Torgerson, 2003; Wertheimer, 2014). Through social media's new advertising opportunities, clinical study conductors have started to utilise social media as a recruitment tool (Darmawan et al., 2020; MD Connect, 2017; Sligh et al., 2016). These new advertising strategies for clinical studies are being implemented by clinical staff as well as professional in-house services and specialised national or international agencies (e.g. *PatientCentra* or *AutoRecruit*).

Recent clinical studies using social media advertising for their recruitment efforts have shown promising results for different parameters and study designs.¹ For an HIV vaccine clinical trial, Sitar et al. (2009) demonstrated that advertising on social media was successful in reaching a wider audience. Frandsen et al. (2014) were able to prove the viability of social media advertising as an addition to other, more traditional recruitment methods for smoking cessation clinical trials. Wisk et al. (2019) found that recruiting college students with type I diabetes to a longitudinal intervention trial via social media was feasible, efficient and

¹Institute of Cognitive Science & Institute of Philosophy, Osnabrück University, Osnabrück, Germany

²Institute of History and Ethics in Medicine, School of Medicine, Technical University of Munich, Munich, Germany

Corresponding author:

Rainer Mühlhoff, Institute of Cognitive Science & Institute of Philosophy, Osnabrück University, Osnabrück 49069, Germany.
Email: rainer.muehlhoff@uni-osnabrueck.de



acceptable, yielding a representative sample of the user-base from which they were drawn. Guthrie et al. (2019) revealed that social media advertising was feasible for clinical trial testing interventions to treat post-menopausal vaginal symptoms. Proponents of these studies stress that social media advertising for participant recruitment has various social benefits, such as increasing ‘fairness’ in access (Whitaker et al., 2017), promoting ‘trust’ and ‘transparency’ (Close et al., 2013; Denecke et al., 2015), or bringing hard-to-reach populations into research (Caplan, Friesen, 2017).

Balancing this often one-sided and positive perspective, we maintain in this article that social media advertising in clinical study recruitment carries a number of ethical and societal risks that have to be assessed and weighed up more thoroughly. The long-held view that social media advertising should be ethically evaluated in the same way as offline strategies (Gelinas et al., 2017) needs to be reconsidered. We highlight the fact that using data on the engagement of single users with clinical ads, social media platforms are able to train machine learning models to predict medical conditions of any third person based on their behavioural data. This poses a range of severe risks, including privacy risks and risks of unequal treatment. With both regulators and study operators as potential audiences in mind, we argue that an unregulated implementation of social media advertising for clinical studies poses significant risks to society and to the quality of research.

We start this article by differentiating four subcategories of ad targeting mechanisms (see ‘What is social media-based advertising and how does it differ from offline ads?’ section and Table 1): static, explicit, predictive and lookalike targeting. This list is ordered according to increasing sophistication in targeting techniques, which also come with growing ethical concerns. After establishing our typology, we will analyse four pressing ethical issues (R1–4, see ‘Risk analysis and ethical evaluation’ section) in relation to social media advertising for clinical trials. These issues are multidimensional, posing potential risks to individual patients, severe risks to society through the aggregation of medical data in the hands of commercial platform companies, and potential threats to the independence and effectiveness of clinical research.

In principle, much of our analysis also holds for targeted advertising outside social media sites, that is, open-domain advertising services such as Google Ads that place targeted ads on any website. We still focus on social media advertising in this article to reduce the complexity of our investigation. Moreover, in our discussion, we use *Facebook* as our main example of a social media platform while pointing out that our analysis holds in an equal way for other social media platforms²

What is social media-based advertising and how does it differ from offline ads?

We will first discuss in what ways social media advertising is different from existing forms of offline advertising for patient recruitment. Social media advertising has common roots with outdoor and print advertising on the streets, in the subway or in magazines. But with recent technological developments, qualitatively new features emerged in online advertising, which enable novel mechanisms of individualised targeting. To highlight these differences, we present a progressive list of four types of targeting mechanisms (Table 1), of which only the first and weakest is comparable to offline mass advertising.

Type A: Static audience targeting

A common form of offline advertising is street or magazine advertising. In patient recruitment, poster campaigns, for instance in subways, print magazines or medical waiting rooms are used to statically target potential participants. In the online realm, static targeting – similarly – uses digital advertising displays, such as banners or text snippets on websites: Advertisers choose a location, for example, a specific website, for a fixed advertisement that is displayed to all visitors of that location over a certain period of time. While the location is carefully chosen to match a specified audience, typically every person that passes by sees that advertisement, and the same version of it. We call this form of targeting ‘static targeting’ because there is no individualised (per visitor) and/or responsive decision mechanism

Table 1. Qualitative comparison of targeting types.

Targeting type	A: Static	B: Explicit	C: Predictive	D: Lookalike
Targeting criterion	Location of the ad display	Personal attributes as provided by users	Personal attributes as predicted by algorithms	Sample cohort of known users
All users generally see the (same) ad?	Yes	No	No	No
Targeting decision based on information the user can control?*	Yes	Yes	No	No
Targeting decision driven by factual criteria?	Yes	Yes	(allegedly, yes)	No

*This implies the certainty that a user does not get targeted by information they did not explicitly provide.

that selects different versions of the advertising for different people walking by.

Static targeting is weak in comparison to more elaborate contemporary online targeting methods and does not play an important role in today's online advertising, as it represents the technological state-of-the-art in the pre-social networking era of the internet (roughly until the turn of the millennium).

The paradigm shift: From offline to online

Indeed, static targeting misses all the specific advertising opportunities afforded by social media. In this case, with 'social media' we broadly refer to digital networked services that implement techniques to uniquely identify their users and their users' real-world relationships (to other users, to brands, affiliations, topics, medical issues, etc.). 'Identification' means that subsequent interactions of the same real person with the service can be matched by the service to belong to one and the same user (represented by a digital 'profile', account or any other kind of persistent identifier), although this kind of identification does not necessarily rely on collecting information about the civic identity of a person. Thus, not only do platforms such as *Facebook*, *Instagram*, *Twitter*, *TikTok*, *Grindr*, *Craigslist*, etc. belong to what we call 'social media', but so do patient platforms like *patients-like-me.com*, or *inspire-im.com*'s patient portal (InspireIM, 2020).

The technological development of social media enabled a profound revolution in advertising due to its ability to aggregate extensive, high-resolution individual data profiles (see e.g. van Dijck, 2009). These data stocks can be leveraged to make automated, fine-grained targeting decisions on a per-user basis that exhibit two largely new features compared to the offline world:

1. Digital media can show different advertisements to different individuals. This could also include different versions of the same advertisement tailored to a user's characteristics or interest profiles.
2. Social media ad providers can collect real-time information about the *performance of an advertisement at user level*. While aggregated information about the performance of an ad could also be measured in offline targeting (e.g. by counting the number of visits to a specific URL), this evaluation can be done at the individual level in social media. This allows the addition of a user's engagement (or disengagement) with an ad to the data profile that is collected about a user. This data, when available on millions of users, can be used to make predictions about which ads perform best with different types of users to inform future advertising decisions.

These two media-technological capacities of online as opposed to offline media enable various qualitative leaps in ad targeting methodology. This will be mapped out as

we now continue the list of increasingly sophisticated targeting principles.

Type B: Explicit targeting

A first qualitative leap occurs when the targeting decision is based on a per-user matching of targeting criteria against explicitly recorded personal attributes. On social media platforms, users sign up with a personal profile and provide personal information through engaging with the platform. In the form of targeting we call 'explicit targeting', the matching operates against data fields that users explicitly enter in their profiles, such as age, gender, relationship status, personal interests, group memberships, friends, educational background, work affiliations, etc.

Providers of explicit targeting are mostly large platform companies because they need a large base of users who are incentivised (through mechanisms often unrelated to advertising) to provide as much personal information to the system as possible. On *Facebook*, for instance, users are motivated to submit personal information to the platform when they communicate with other users (likes, photos, messages, ...). Monetising this information through advertising is a secondary use, albeit the main source of revenue for the platform (Statista, 2022).

Compared to static targeting, explicit targeting frequently renders the specification of a *location* (e.g. a specific website) obsolete. The location, to begin with, was always only a proxy for personal attributes, which can be explicitly specified in this new form of targeting. Predetermined by the individualised targeting criteria, the platform automatically issues at which location (i.e. at which point in the interaction between user and online service) which ad is to be shown.

In the case of patient recruitment for clinical studies, the explicit targeting method shifts attention from determining the best physical location of an advertisement to naming the key demographic and personal characteristics of the target audience. For example on *Facebook*, it is an easy task to target an advertisement to single female users from Germany, aged 18–65 years who are interested in 'Smoking or Nicotine replacement therapy' (Figure 1). While explicit targeting thus promises to be specific to key demographic requirements for the cohort of a clinical study, it still provides a comparatively coarse targeting mechanism, which suggests that explicit targeting is only affordable to clinical studies that recruit for very broad audiences and large study populations.

Type C: Predictive targeting

The next qualitative leap comes with individual targeting mechanisms that not only rely on personal attributes explicitly specified by the users, but also on predicted

The screenshot displays the Facebook ad configuration interface. On the left, there are three main sections: 'Special Ad Category' with a toggle switch, 'Audience' with three radio button options, and 'Audience Details' with a 'Create New' button. On the right, there are sections for 'Estimated Daily Results' and 'Payment Summary'. At the bottom, there is a 'Promote Now' button and a disclaimer.

Special Ad Category [Learn More](#)
 Ads about social issues (such as the economy, or civil and social rights), elections, or political figures or campaigns

Audience
 Who should see your ad?

People you choose through targeting

Example Predictive Targeting

Example Explicit Targeting

Audience Details

Location - Living In Germany
 Age 18 - 65+
 People Who Match Interests: Smoking or Nicotine replacement therapy

Create New

See All Previews

Estimated Daily Results

People Reached ⓘ 93 - 270
 Page Likes ⓘ 0 - 8

Payment Summary
 Your ad will run for 5 days.

Total budget €10.00 EUR
 €2.00 a day x 5 days.

By clicking Promote Now, you agree to Facebook's [Terms & Conditions](#) | [Help Center](#) **Promote Now**

Figure 1. Screenshot Facebook.com: Configuration of an explicitly targeted online ad.

information. Imagine a user did not specify their age or gender in their profile information, while their extensive behavioural data (posts, likes, browsing history or social graphs) are available to the platform. Studies showed that the missing attributes can be estimated from the available data (Kosinski et al., 2013). Using advanced machine learning and data analytics techniques, what is referred to as ‘predictive targeting’ therefore leverages behavioural, usage and tracking data to estimate relevant targeting attributes about single users where such attributes are not explicitly given.³

This method, however, is not limited to ‘filling the gaps’ in users’ profile information. Rather, this technique enables a covert expansion of the available targeting parameters to include information that hardly any user would voluntarily provide and whose explicit processing might even be prohibited by data protection regulations. This information includes, for example, purchasing power, wealth or credit risk scores, as well as estimated information on medical conditions, sexual orientation, ethnic or racial background (see e.g. Mühlhoff, 2021; O’Neil, 2016). As Andreou et al. (2018) explain, *any* data that users share with the

platform is treated as ‘raw data’ input that can potentially be ‘translated’ by data analytics algorithms into ‘targeting attributes [...] that advertisers can specify to select different groups of users’ (Andreou et al., 2018).

This step of ‘translation’ typically comes with an increase in the sensitivity of information, or infringement of privacy, as Mühlhoff (2021) argues. ‘Predictive analytics’ enables advertising providers to estimate highly sensitive and intimate information about users from less sensitive and often readily available data (cf. Mühlhoff, 2021; Figure 1). Among the information that is predictable from social media usage data are many medical data fields that might be of particular interest to advertisers in the field of patient recruitment for clinical studies. For instance, researchers at the University of Pennsylvania have shown that diseases such as depression, psychosis, diabetes or high blood pressure can be predicted from a user’s postings on *Facebook* (Merchant et al., 2019). Kosinski et al. (2013) showed that sexual identity could be predicted from the same data.

It is a reasonable assumption that most contemporary online advertising providers combine explicit and predictive

targeting (types B and C) without allowing advertising customers to opt out of the predictive mechanism.⁴ Advertising customers just specify certain targeting criteria that will be matched by the social media platform either with data explicitly provided by users (where available, e.g., for users who disclose their age in their profile) or with predictions of the same parameters when a user did not provide that information about themselves. In social media advertising for clinical studies, this would mean that medical – and thus potentially sensitive – attributes of users are being estimated based on data from other users that disclose such information. This could range from alcohol intake and smoking habits to diabetes or infectious diseases such as human immunodeficiency virus (HIV), hepatitis B or other highly stigmatised indications.

Type D: Lookalike targeting

Predictive targeting, as an ‘improved’ version of explicit targeting, makes ‘virtual’ (or estimated) audience attributes available as targeting criteria, in addition to the attributes users would specify about themselves. The next step in the increasing development of targeting methods presents a way to do without the specification of personal attributes as targeting criteria by the advertiser. In what we call ‘lookalike targeting’, advertisers provide a list of relevant *users* who are known to respond positively to the campaign. A targeting algorithm is then used to find *similar* users on the platform (with the criteria for similarity being determined ‘algorithmically’), who are then shown the advertisement. In our terminology of ‘lookalike’ targeting, we follow *Facebook*, which introduced so-called ‘lookalike audiences’ in 2013 (cf. Figure 2). Advertising customers can specify a list of users (e.g. a few hundred) and order *Facebook* to display a certain advertisement to those users who ‘look like’ the sample. This means that data analytics and machine-learning algorithms automatically determine behavioural and demographic markers that are specific to the sample group and then apply this analysis to all other *Facebook* users to show the advertisement to those that are deemed similar (Andreou et al., 2019; Semerádová, Weinlich, 2019).⁵

The power of this targeting principle becomes evident, for example, in combination with *Facebook* groups, fan pages or the collection of email addresses from offline events. *Facebook* groups are online interest groups where users subscribe in order to engage with other users about a shared interest, such as arts and entertainment, products, leisure activities, political views or a shared disease or medical condition. Lookalike audiences can be created from the members of such groups as well as from the followers of political organisations, activist groups, religious associations or minority movements. Alternatively, advertisers can upload lists of email addresses of known customers (or patients, in the case of patient recruitment, see Ryerson,

2021a) from which *Facebook* creates a lookalike audience. This targeting type is known to have been mobilised in the Donald Trump presidential campaign in 2016 (cf. Baldwin-Philippi, 2017). In commercial patient recruitment for clinical studies, *Facebook*’s lookalike audience feature has been recommended as an ‘effective’ tool for online micro-targeting, for example, by *antidote*, a patient recruitment agency that advertises ‘precision recruitment services’ for clinical trial sponsors on their website (Ryerson, 2021b).

The hidden epistemology of online targeting

The broader transformation from static to lookalike targeting marks a fundamental change in the implicit epistemology of advertising. The first transition, from static to explicit targeting, suggests a return to a positivistic reliance on explicit information about the single individual (cf. Kitchin, 2014). Whereas static targeting used locations (e.g. street corners, metro lines, magazines or event locations) as proxies to hit a certain market segment in a time when little direct individual data was available, explicit targeting makes use of the new wealth of information people have been disclosing about themselves on social media networks since the 2000s. This can be seen as the beginning of the era of ‘micro-targeting’ (Barbu, 2014).

This new positivism, however, outdoes itself when it transitions from explicit to predictive targeting. Here, a speculative element creeps into the underlying data epistemology because predicted attributes are often derived as the ‘best bet’ from statistical inferences (Joque, 2022; Mühlhoff, 2021). This represents a shift from relying on how people reflect on themselves to how they behave. The kind of reasoning that is implemented by these algorithms remains positivistic, but turns from subjectivism to the *digital behaviourism* of pattern matching. Remarkably, this shift also breaks with the individualism inherent to subjective information. Instead, it uses *relational* information about how users compare to each other to derive supposed similarities in *personal attributes* from similarities in *behaviour* (Mühlhoff, 2021).

Given this relational approach of predictive behavioural epistemologies, it seems logical that in the next step of the development of targeting methods, leading from predictive to lookalike targeting, we see the use of relational information coming fully to the fore. Lookalike targeting radicalises this approach by offering an alternative to specifying a list of personal attributes for ad targeting (which still makes the advertiser focus on the individual): advertisers simply specify a list of known users. This list fixates a relational constellation of existing persons instead of a grid of individual attributes, so that a matching algorithm can establish who else is in close proximity to those users.

In the context of patient recruitment in clinical studies, it is important to note that determining the similarity metric (i.e. the behavioural markers that determine the relevance

Figure 2. Screenshot Facebook.com: Creation of a lookalike target audience, showing the different attributes that can be specified, for example, email, phone number, names, gender, age, etc.

of a person to the clinical study) is largely delegated to secret algorithms and their designers when proprietary targeting services are used. Specifying relevant personal attributes, seen in explicit and predictive targeting, is a way of establishing factual criteria for the potential relevance of a given person to the clinical study. This relation between criteria and advertising is, at least in principle, accessible to scientific explanation (e.g. based on knowledge about the prevalence of a certain disease in a demographic cohort or in relation to certain lifestyle aspects). Lookalike targeting breaks with this chain of reasoning. Ads will be shown to a person X not because they belong to a cohort that is statistically likely to suffer from a certain medical condition, but because an opaque,

proprietary algorithmic procedure deems that it is profitable to show the ad to X (Joque, 2022).

In the next section we will discuss how the various stages of this epistemological transformation come with different ethical, data protection and epistemological risks.

Risk analysis and ethical evaluation

Multiple ethical and data protection concerns arise in relation to advertising clinical studies on social media. In this section, we put forward a selection of such concerns (risks R1–4, Table 2), differentiating according to the typology of targeting methods laid out in the previous section.

Table 2. Risk comparison of targeting types.

Targeting type	A: Static	B: Explicit	C: Predictive	D: Lookalike
R1: Violations of privacy	No	No Targeting uses information the user explicitly disclosed	Yes Violation of users' 'predictive privacy'	Yes Violation of users' 'predictive privacy'
R2: Disclosure of medical information to the platform	(yes) (if online and if ad engagement is measured and linked to user-level behavioral data)	Yes If ad engagement is measured and linked to user-level behavioral data	Yes Training the targeting model is approximately equivalent to training a predictive model of medical information	Yes Training the targeting model is approximately equivalent to training a predictive model of medical information
R3: Unethical exploitation of the limits of 'informed consent'	No	(no) We assume that users are properly informed about the use of their data for advertising purposes	Yes Sensitive or personal information is being predicted about users without their consent	Yes Behavioural resemblances of users are calculated without their consent
R4: Impairment of research quality	No	No	Yes Black box algorithm interpolates known attributes about users for participant selection	Yes Black box algorithm selects users who 'look similar'

For every targeting type (columns A–D) it is indicated in what way the different risks (rows R1–4) are prevalent.

R1: Violations of individual privacy

Social media targeting algorithms match audience information provided by the advertising customer with social media users. As we saw in the previous section, state-of-the-art predictive targeting mechanisms (type C) do not require that all users explicitly specify the attributes in question, as missing or undisclosed information is routinely estimated by predictive models. Moreover, predictive targeting technologies extend the range of available targeting attributes to personal information that users typically do not enter themselves.

The first ethical concern connected to this technology is that it violates users' privacy in a new way. Under the title of 'predictive privacy', Mühlhoff (2021) pointed out that privacy can be violated by *predicted* (instead of only recorded) information. Following that approach, we define the 'predictive privacy' of an individual or group to be violated:

'if sensitive information about that person or group is predicted against their will or without their knowledge on the basis of data of many other individuals, provided that these predictions [could] lead to decisions that affect someone's social, economic, psychological, physical, ... well-being or freedom'. (cf. Mühlhoff, 2021: 5)

To see in which way this definition applies to the present scenario, we have to examine three conditions. First, predictive targeting generates information about users

without their knowledge and potentially against their will. As information that is predicted for targeting of clinical studies includes, or correlates with, sensitive information such as addictions, health, psychological conditions, sexual orientation, etc., we must assume that the application of predictive targeting in this context is very likely to involve information users would not willingly provide to the platform.

Secondly, predictive and lookalike targeting classify users through large-scale comparisons with the behavioural data of many other individuals and are thus enabled by the data that social media users *collectively* generate. Data analytics techniques 'learn' behavioural markers of the sensitive target attributes by pattern recognition in seemingly less sensitive and readily available data, such as tracking or usage data, that are collected from all users of the respective platform. O'Neil (2016) characterises this filtering mechanism as a 'people like you' principle (p. 145). If the system learns that a group of users who share a certain sensitive attribute (e.g. 'parental separation') shows a significant correlation in behavioural data fields (e.g. in their way of using a social media platform), this analysis is then used in a reverse direction: if a new person correlates with that group in terms of behavioural data, it is then concluded that they likely also share the sensitive target attribute.⁶

The technological possibility of predictive targeting is an effect of collective data aggregation, to which most of us regularly contribute through our use of networked platforms. The feasibility of this technology is unaffected by

any single user's decision to not disclose a certain piece of information, as long as there is still a large cohort sharing much of their data (and even if they do so anonymously). Predictive targeting therefore poses a challenge to all individualistic and control-oriented interpretations of privacy rights as implemented, for instance, in 'informed consent' policies (Blanke, 2020; Kröger et al., 2021; Wachter, 2019). As predictive algorithms can fill the data gaps that might result from single users' deliberate non-disclosure of certain information, these users can still be targeted. By not joining the game, single users can neither prevent predictive analytics being applied to others nor to themselves. This is counter-intuitive and deceptive to many users who are not familiar with the predictive targeting methodology. A lack of public awareness thus contributes to the largely unregulated applicability of this kind of data processing.

Thirdly, to meet the requirements for 'violation of predictive privacy' according to the definition from above, we need to establish in what way the prediction of hidden targeting attributes could affect 'someone's social, economic, psychological, physical, ... well-being or freedom'. At this point, we need to consider the type of information that is implicitly or explicitly predicted in our concrete cases. In the context of clinical studies, reaching effective targeting decisions often correlates with an assessment of medical information about social media users since the advertising seeks to reach persons with specific diseases or medical (pre-)conditions. This makes training an effective predictive model for this ad targeting decision *equivalent to building a predictive model of medical information*. As medical information is classified as particularly sensitive by many standing data protection regulations such as the General Data Protection Regulation (GDPR) in the EU and the Health Insurance Portability and Accountability Act (HIPAA) in the United States, the violations of predictive privacy that are facilitated by targeting algorithms for clinical studies are particularly severe.

In what way, though, could this information be used to anyone's detriment? Tufekci pointed out that 'when 'big data' reveals more than was disclosed', these 'computational violations of privacy' have potentially severe 'civil rights implications' (Tufekci, 2015: 209). O'Neil makes clear that predictive models are being used in a wide range of industrial applications with material consequences for individuals, ranging from (semi-)automated hiring processes on the job market to insurance and credit scoring to the criminal justice system (O'Neil, 2016). As O'Neil argues, the way predictive models are being used is 'exacerbating inequality and punishing the poor' (p. 17). Behind these analyses of the consequences of predicted information is the observation that predicted information is generally used for a differential treatment of users.

The urgency of an ethical and political debate of predictive violations of privacy is further underpinned by the

observation that despite the extra sensitive status granted to medical information by many data protection regulations such as the EU's GDPR, most of these regulations *fail* to effectively protect against building and utilising predictive models.⁷ It has been pointed out that the GDPR is ineffective in protecting against the use of predicted information because it tends to regulate data processing at the input stage rather than the use of derived information (Wachter, 2019; Wachter, Mittelstadt, 2018); its provisions against profiling are limited to fully automated systems (Art. 22 GDPR) and most social media platforms rely on user consent as the legal ground for the processing of their data, something which is easy for the platforms to obtain (Mühlhoff and Ruschmeier, 2022).

R2: Collective privacy implications: Disclosure of medical information to the platform

An even stronger ethical concern regarding the violation of predictive privacy arises when we do not focus on the individuals that are part of a clinical study or that see the advertisement for a study on social media, but shift our attention to the effects on 'digital bystanders' as well as simply *anybody* else on the same social media platform.

To show why this is a concern, our main argument in this article is that social media advertising for clinical studies *enables social media platforms to train specialised predictive models of medical conditions* in *any* of the (present and future) platform users. Online advertising platforms measure user engagement with the ads they show. In the case of advertising for clinical studies, this engagement data reveals medically relevant information, which the platform collects in addition to the data it already has about the user. The platform uses this combined dataset as training data for the targeting algorithm, at the core of which is a predictive model that estimates the likelihood of a user to click on that ad. In case of medical ads, as it was argued in R1 above, such a predictive model is effectively an engine to predict *medical information* about *any* platform user – now or in the future (cf. Figure 3).

While the exact workings of targeting algorithms are trade secrets, it is not only theoretically plausible but also empirically evident that social media platforms indeed leverage the ad engagement data to make their models 'learn'. For instance, a 2018 study by Borodovsky et al. (2018) made use of *Facebook* advertising to study cannabis use behaviours in the United States. As the authors describe, while running *Facebook* ad campaigns to find cannabis users they observed that targeting effectiveness increased after a short learning period:

'[T]he algorithms used to distribute advertisements appear to require a sufficient amount of time to complete a learning phase before they become effective. During this learning

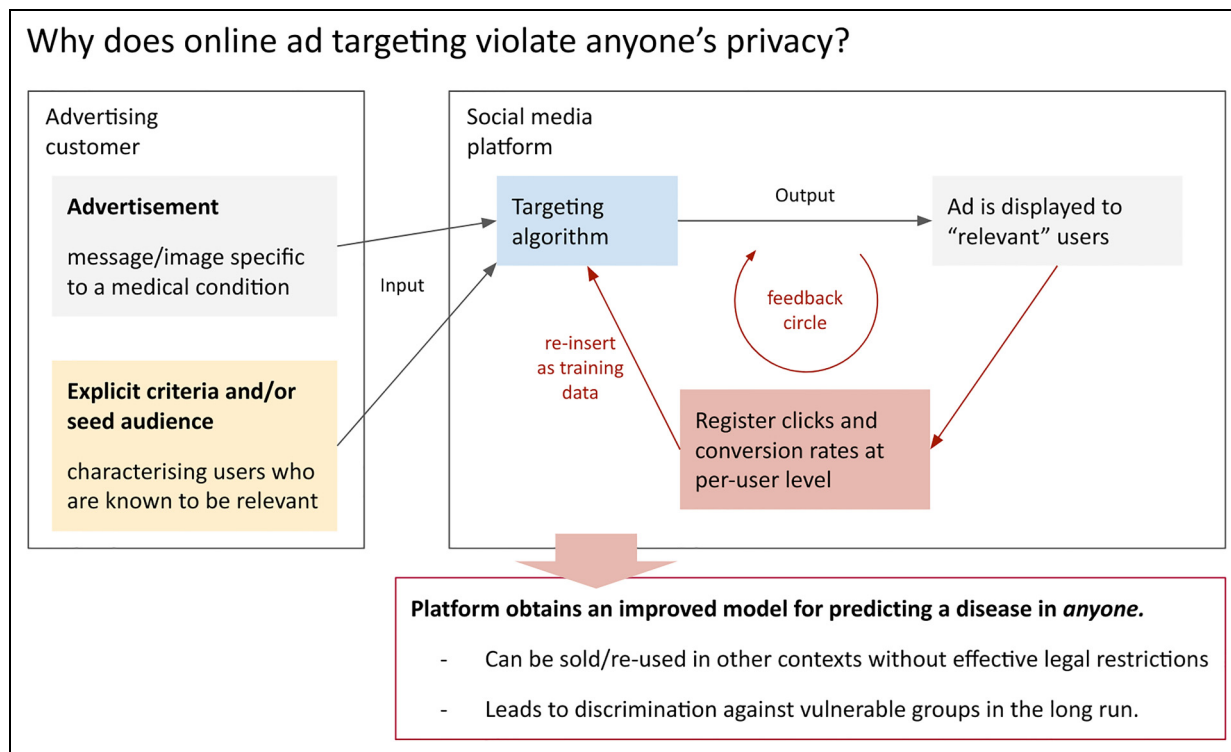


Figure 3. Advertisers help data companies to train predictive models.

phase, enough data must be accumulated to determine which members of the target population have the highest probability of engaging with the advertisement. [...] Currently, when we use Facebook advertisements to reach cannabis users, we begin advertising with a low spending limit of US \$10 per day for 48 to 72 h, which we believe facilitates algorithm learning. After this 48- to 72-h time frame, we have consistently seen a notable increase in the rate of clicks'. (Borodovsky et al., 2018: 2)

Given the many opportunities of re-utilising these targeting models as predictive models of medical information and related risks, the social media platform benefits twice from the advertising deal. While study operators pay for access to the results of the predictive model, platform companies can re-use *the predictive model* in other sectors such as risk assessment in insurance and financing. The clinical study operator, which is often a publicly funded institution, thus substantially (but probably unintentionally) assists the platform in building a predictive model of specific medical information, the further use of which it cannot control. The trained model is then proprietary data owned by the platform company. The trained model's internal parameters do not generally classify as sensitive or personal data according to existing data protection regulation, such as the EU's GDPR. Thus, the model can be repurposed by the platform for any other objectives or even be sold to

third parties, instead of being at the disposal of medical research and services.

In the hands of private actors in a largely unregulated market, such predictive models pose an immense threat to social justice and equality of treatment. They endow platform companies with technological capacities to treat individuals and groups differently at large scale, which marks a structural power asymmetry between society and business interests. As Barocas and Selbst (2016) argue, this power asymmetry exacerbates already existing structures of inequality and unfair discriminations when predictive algorithms 'discover surprisingly useful regularities that are really just pre-existing patterns of exclusion and inequality' in our societies. This is why an 'unthinking reliance' on, and unregulated spread of those technologies 'can deny historically disadvantaged and vulnerable groups full participation in society' (Barocas, Selbst, 2016: 671).

Consequently, the moral responsibility of study operators is not inconsiderable. Note that even a *single* decision by study operators to use social media advertising can have a significant impact on society, as even a single ad campaign can produce a well-trained predictive model for a specific medical condition. Admittedly, the platform companies are the primary actors in this context, because they can undertake or refrain from the production of corresponding predictive models. In addition, big data and artificial intelligence (AI) businesses that rely on the use of (potentially anonymised) aggregate data are currently heavily

unregulated, which is a major problem that needs urgent political attention (cf. Mühlhoff, Ruschemeier, 2022). Nevertheless, until better regulation is in place, clinical study operators also bear a heavy responsibility. What is at stake is that private companies are enabled, with the help of public research, to discriminate against already vulnerable user groups along estimated medical parameters. It needs to be highlighted that the data protection risk outlined here is a genuinely collectivist issue. This means that the data protection claim cannot be reduced to individual claims and the responsibility cannot easily be pinpointed to one actor. The case of advertising for clinical studies shows that better political regulation of the entire field of predictive analytics and its actors is required.

It should also be mentioned that the data protection risk outlined in R1 and R2 is not restricted to the predictive forms of targeting (types C & D). It pertains to all digital advertising services that collect user-level engagement data (e.g. click rates).

R3: Exploiting the limits of (biomedical) research ethics

Protection of participants' privacy and data protection rights have been well-established as a central pillar of biomedical research ethics through international frameworks and guidelines, such as those established by the World Health Organisation (WHO) and the Belmont report (CIOMS, 2016; The World Medical Association, 2013; cf. Beauchamp, Childress, 2012). According to the WHO, the collection, storage and use of data in health-related research must be governed by an informed consent process similar to those related to participation in the research procedure itself (CIOMS, 2016). Researchers, according to the WHO, are obliged to 'anticipate, control, monitor and review interactions with their data across all stages of the research' (CIOMS, 2016: 83). In this subsection, we argue that the WHO provision also applies to the participant recruitment stage.

In light of the previous discussion, we suggest that the frameworks be extended or sharpened corresponding to the specific challenges of social media recruitment strategies and technologies. There are at least *three* distinct data ethical challenges originating from social media advertising that should explicitly be dealt with in research ethics guidelines:

1. Consent to be contacted: Online targeting involves implicit or explicit estimations of medical information about a large group of social media users before they even see an advertisement. If social media advertising is utilised for study recruitment for stigmatising or rare diseases, particularly sensitive medical information is estimated and processed by the platform *before* a

contact between the potential candidate and the study operator is established, and thus, before the study operator could ask for consent. This preselection of users is in direct conflict with the obligation to obtain informed consent from any potential participant before contacting them for recruitment. As Arigo et al. (2018) argue, even if explicit consent were obtained, for example, together with other terms and conditions that users are asked to agree to, the current standard practices of having users to agree to such conditions by ticking a box might be largely ineffective for user information, as users are too used to checking off terms without reading them carefully, if at all.

2. Outsourcing 'dirty business': Existing biomedical research ethics guidelines are not devised for social media targeting and can easily be circumvented. For instance, guideline 22 in the WHO framework (CIOMS, 2016) demands that: 'Researchers collecting data on individuals and groups through publicly accessible websites without direct interaction with persons should, at a minimum, obtain permission from website owners, post a notice of research intent, and ensure compliance with published terms of website use'. (CIOMS, 2016: 83)

Due to the architectural differences of social media platforms compared to other publicly accessible websites to which the guideline refers, it is – in the case of social media advertising for clinical studies – almost impossible to comply with the requirements mentioned therein. Using a social media advertising service outsources a tremendous amount of data processing from the study operator to the social media platform. As a result, researchers (unwittingly) bypass severe ethical obligations that actually rest with the study operators. Moreover, while social media users give their *general* consent to data processing by the platform, they do not explicitly consent to the concrete advertising campaign and the big data side-effects of a given clinical study advertisement. As argued above, many users are unaware that their data is being used for these purposes.

With this in mind, the use of social media platforms for advertising clinical studies presents an outsourcing of the ethically doubtful data processing work of targeting to the platform company. This opens up the possibility for study operators to make use of the advanced predictive technologies of social media platforms and have their data processed in ways they would ethically, legally and practically not be able to do themselves. To protect patients as well as clinical study operators, the social media-specific ethical considerations about data collection and processing should be covered in ethics guidelines (Hunter et al., 2018). Moreover, entirely new forms of informed consent procedures are required for specific cases that are unique to social media, for example, when researchers want to

recruit in the networks of their study participants (Custers et al., 2014; Gelinas et al., 2017; O'Neil, Schutt, 2013).

1. With predictive or lookalike targeting mechanisms, a user's health is predicted on the basis of characteristics they did not explicitly disclose. This is likely to cause harm because individuals are confronted with a (potential) diagnosis without having chosen to be so (cf. Thornton et al., 2016). As discussed above (see R2), the feedback mechanisms of online ad targeting allow the platform to train a *general* model to predict health conditions of *any* user, including those who do not participate in the study and those who never see an ad or otherwise interact with the campaign or research. Again, provisions such as the WHO framework (Council for International Organizations of Medical Sciences, World Health Organization, 2016) and the Belmont report (The World Medical Association, 2013) were not written with this implication in mind and therefore need updating for the societal-scale data externalities that derive from individual data processing.

Arigo et al. (2018) second this collectivist line of argumentation by pointing out that it needs to be considered how the infringement of privacy might affect not only potential patients, but also 'bystanders' such as relatives or other people with social (media) connections to potential participants. We go even further in pointing out that this issue shows a fundamental limitation of the informed consent principle which ignores that the decision made does not have consequences only for the person asked, but for society at large (cf. Kröger et al., 2021; Mühlhoff, 2021). Ethics guidelines, which reduce data protection to informed consent, therefore not only neglect one of the most serious challenges in data ethics and data protection in the context of social media, but also *endorse* harmful data processing under the flag of ethical and GDPR compliance. As a result, most existing guidelines are proving to be toothless or even counter-productive in regard to the potentially severe societal effects of contemporary big data analytics.

R4: Impairment of the quality of (biomedical) research

Our analysis shows that delegating ad targeting decisions to Big Data-based advertising services buys into a comprehensive amount of 'magic' in deciding who is, and who is not, a potential candidate for a study. From the point of view of the researchers operating a clinical study, this means outsourcing targeting decisions to a 'black box' algorithm (cf. Pasquale, 2016). In this subsection, we argue that this has potential implications on the quality of the research in question.

The actual algorithmic functions of targeting mechanisms are typically not disclosed by social media platforms.

Researchers working outside the respective platform, including recruiting researchers, usually have no control and often no knowledge of which data is collected, stored and used in the targeting process. Neither can external researchers control or ensure whether the algorithms have been updated during the recruitment phases by the respective platform and the targeting mechanisms have therefore undergone changes (Hitlin et al., 2019). These circumstances create a methodological pitfall according to existing ethics frameworks, which demand that study conductors publish not only the results of their research on human subjects, but also account for the completeness and accuracy of their reports, including disclosure of recruitment strategies and cohort emergence (The World Medical Association, 2013). As the CIOMS framework highlights:

'groups, communities and individuals invited to participate in research must be selected for scientific reasons and not because they are easy to recruit because of their compromised social or economic position or their ease of manipulation'. (CIOMS, 2016: 7)

Social media-based recruitment violates these principles and thus sacrifices indispensable methodological rigour and control to commercial owners of data and algorithms. As recruiting decisions are central to a sound study design (see e.g. Patel et al., 2003; Puffer, Torgerson, 2003), adoption of social media-based recruitment strategies outsources to commercial actors what is actually *the researchers'* work. Given the architecture of their business models, social media platforms are not in a position to provide the same standards of scientific and methodological transparency and accountability. Thus, if used extensively, social media advertising poses a threat to (biomedical) research quality. This threat includes, but is not limited to, the uncontrolled potential of skews in the study due to uncontrolled biases in the recruited cohort. This could affect conclusions drawn from the study and thus patients who later receive treatment on the basis of the study results.⁸ This issue also relates to the principle of justice in biomedical research ethics, which forms the basis of a key element of health policy: the equitable sharing of healthcare resources (Beauchamp and Childress, 2012).

In fact, these circumstances described above lead to a seemingly paradoxical constellation: social media recruitment is often advertised in the name of justice for its acclaimed capacity to reach poorly accessible and vulnerable patient groups that are otherwise hard to reach (Bonevski et al., 2014; King et al., 2014; Wozney et al., 2019). However, studies for which this argument applies are usually characterised by inclusion or exclusion criteria based on particularly sensitive medical information. Hence, deploying social media advertising in this situation could supposedly target more individuals of a highly specific medical situation, but at the cost of an increased

opacity and unaccountability in the targeting mechanisms, which compromises scientific rigour. In such cases, we also have to refer back to the increased severity of the predictive privacy implications (R2). Hence, the issues of justice and accessibility are easily turned around, with predictive targeting ultimately working against those groups who are supposedly better included by such measures.

Putting it all together: Austerity in the guise of benevolence

Social media-based advertising is often touted as *the* step forward in recruiting for clinical studies in the context of rare or stigmatised diseases. Scholars suggested that social media can help reach the poorly accessible and often vulnerable populations, for instance, persons suffering from specific mental illnesses, addictions, or sexual and ethnic minority groups (Caplan, Friesen, 2017; Sedrak et al., 2016; Wozney et al., 2019). While this claim of a greater recipient reach may be true from a limited short-term perspective, from a structural and societal viewpoint, our ethical assessment casts serious doubts. It should be emphasised that the purported argument about the quality of advertising measures is in fact only a financial argument. Clinical study operators are usually under immense pressure to optimise for recruitment accrual given a limited budget. Social media advertising seems attractive in this situation as it promises a higher advertising return per cost, especially when targeting scarce interest cohorts and vulnerable groups (Frandsen et al., 2016; Topolovec-Vranic, Natarajan, 2016). Our analysis of risk R2 shows that this seemingly simple way of raising efficiency comes at potentially high societal and individual costs in the long run. From an overarching and societal perspective, these methods run a high risk of making vulnerable, especially stigmatised, patient groups pay for being reached by medical researchers, as their inclusion comes at the cost of greater vulnerability to future discrimination through Big Data and AI-based decision making everywhere in the economy.

The proprietary data aggregated by platform companies as a corollary to the medical ad campaign allows those companies to build AI models that can predict the relevant medical condition in anybody. Those models, once they exist, are set to play a role in the mechanisms of structural power in our societies and will likely not play out to the benefit of the vulnerable groups (Mühlhoff, 2020; O'Neil, 2016). Creation and circulation of those models is not effectively regulated by current data protection legislation such as the EU's GDPR (Wachter, 2019). As both the training data (which is generated by the advertising feedback mechanisms) and the internal model parameters can be handled as anonymised data, data protection regulations do not apply (Mühlhoff, Ruschemeier, 2022). Thus, the collected data can be re-used in scenarios such

as (semi-)automated job selection, insurance pricing, political persecution, access to financial, educational and other resources.

The supposed increase in efficiency of current ad campaigns will thus backfire many times in the form of technologically driven social selection and discrimination. From the perspective of platform companies, such advertising deals are just welcome opportunities to collect user-level data which will be used as AI training data in an epistemic field that is difficult to access. In a context where predictive analytics is as little regulated as it is today, study operators and funding bodies opting to enable or pursue this form of advertising for clinical studies are (unwittingly) complicit in enabling commercially driven forms of future injustice.

These considerations profoundly demystify the promise of benevolence that is often connected to the use of social media advertising for clinical studies. From a more structural and long-term perspective, the real incentive for using social media advertising seems first of all driven by austerity – with the same budget, advertisers seek to reach a more specific audience. *This* efficiency argument, in its short-sightedness, might be true, although the impairment of biomedical research quality stands against this (see R4). Yet, in light of our full-fledged ethical discussion, the answer seems simple: make more resources available for this kind of research. The resources we seek to save now by outsourcing cohort selection to Big Data companies will be paid later (and multiple times) when our societies have to compensate for increased social injustice and societal chasms resulting from AI-based mechanisms of social selection.

At this point, there is a high demand for better regulation. Big Data companies should be effectively regulated in their processing and re-purposing of aggregated and anonymised data in the form of predictive models that can be transferred to other contexts (Blanke, 2020; cf. Mühlhoff, 2021). Clinical research, on the other hand, must be prohibited from implicitly contributing to private companies collecting proprietary data that will harm our societies. At this point, a key to better and more effective regulation is that both researchers and regulators pay more attention to aspects of data protection that go beyond individual privacy. Protecting the *individual privacy of study participants* is not enough to ask for in an ethical evaluation of clinical studies. Even in anonymised form, the data of actual or potential patients is being leveraged by big data companies to learn medical information about *any* other user on their platforms. In this situation, biomedical research ethics guidelines that frantically uphold protection of *individual* privacy, for instance, through anonymisation or informed consent, are even counter-productive as they *enable* operators to pursue damaging Big Data business models under the flag of ethical compliance.

Conclusion

We have unpacked the apparently benevolent usage of social media advertising in the context of recruiting participants for clinical studies. After other studies had examined and highlighted the positive effects of social media recruitment (Frandsen et al., 2014; Guthrie et al., 2019; cf. Sitar et al., 2009; Wisk et al., 2019), we found severe potential for a range of negative effects regarding privacy, social inequality and scientific research quality.

Setting the stage for our investigation, we separated different types of ad targeting. We showed the general threats of online audience targeting (type A), supported by the explanation of the paradigm shift from offline to online. An increasing list of threats was introduced that originate from explicit targeting (type B), predictive targeting (type C) and lookalike targeting (type D).

In our risk analysis and ethical evaluation, we identified different risks originating from social media advertising. Risk R1 starts from the concern that the effect of the established principles of ‘informed consent’ and ‘privacy self-management’ are limited in online spaces (Kröger et al., 2021; Mühlhoff, Ruschemeier, 2022). Regarding individual privacy rights, predictive and lookalike targeting mechanisms are problematic as they derive potentially sensitive (medical) information from less sensitive information about arbitrary users – including those who are not targeted by the ad campaign. This constitutes a violation of users’ *predictive privacy* (Mühlhoff, 2021) that cannot be prevented by informed consent policies.

In risk R2, we transferred this argument to a structural level. We made the case that social media advertising for clinical studies enables social media platforms to build AI models that can predict medical conditions about *arbitrary* platform users. There is no provision in place to prohibit platform companies from re-using the data they glean from an ad campaign (user-level ad feedback) to train predictive models that can be applied in other contexts such as insurance or job recruitment industries. This is a structural and not (only) an individual ethical concern, as potential applications of these models might lead to differential treatment of and discrimination against already vulnerable groups – potentially years after the ad campaign. This showed how the ‘benevolent’ idea of using social media targeting to reach vulnerable patient groups more effectively with the same advertising budget might backfire in the long run by an increase in discrimination and structural injustice against these groups through automated decision making (Mühlhoff, 2020, 2021).

In risk R3 we then analysed the detrimental effects of biomedical research ethics guidelines that focus on privacy as merely an *individual* right as opposed to the structural viewpoint of data protection that includes the societal externalities of aggregated data. In light of the specific challenges of Big Data and AI, we argued that

biomedical research ethics frameworks need updating in order to avoid unwittingly facilitating industry business models that seek to leverage Big Data inferences under the flag of ethical compliance.

In risk R4, finally, we pointed at the potential impairment of biomedical research quality that results from outsourcing recruitment decisions to ‘black box’ targeting algorithms of commercial actors. Patient selection is not scientifically accountable at the same level of rigour if delegated to closed-source social media platforms and their ad targeting algorithms.

We conclude that social media advertising should not be eligible as a recruitment tool for clinical studies so long as the processing of even anonymised social media usage data and the training of predictive models by data analytics and AI companies is not sufficiently regulated. The more specific and vulnerable the target group of a clinical study (e.g. rare or stigmatised diseases) is, the more severe are the structural ethical considerations proposed in our analysis. Moreover, it is likely (but will have to be detailed in future research) that many of the risks we pointed out also hold for targeted advertising services outside social media platforms, such as *Google Ads*, which places ads on any website. We also suggest that the findings from our analysis are transferable to other sectors of medical and behavioural research, and, in fact, of advertising in general. In particular, risks R1 and R2 hold nearly everywhere: Predictive models can be built from advertising feedback data that is routinely and legally collected from online advertising for whatever topic. Finally, we urge study operators as well as ethics boards to factor in the various potential negative effects of social media advertising when examining recruitment options. Overall, we advocate an increase in regulation of the entire social media advertising sector and of predictive analytics technology in general.

Acknowledgments

We thank Anastasija Kocić, Mareike Lisker and Jan Siebold for critical discussions, background research and proof reading of this manuscript at its various stages.

Declaration of conflicting interests


The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Theresa Willem’s contribution to this work was supported by the H2020 European Research Council (grant number 848223).

ORCID iDs

Rainer Mühlhoff  <https://orcid.org/0000-0002-3936-9919>

Theresa Willem  <https://orcid.org/0000-0001-7643-8816>

Notes

- 1 Existing research as well as our discussion largely focus on paid ads. Other, more implicit social media-based advertising strategies exist, such as postings in interest groups, ads shared by influential users or by former participants. While we ignore these other ways to focus our analysis, it is safe to assume that most of the risks also pertain to more implicit strategies. Figuring this out in detail should be a matter of follow-up research.
- 2 When referring to *Facebook* throughout this article, we reference the platform, not *Meta*, the company.
- 3 In critical debates, the principle of predictive targeting was often discussed in the context of predictive policing and predictive criminal prosecution, cf. for example, Ferguson (2016). An important precursor to our debate is also the critique of ‘predictive profiling’ by Hildebrandt and colleagues, cf. Hildebrandt and Gutwirth (2008) for an extensive analysis.
- 4 Some platforms, including Facebook, do allow their *users* to opt out from seeing predictively targeted ads. Provided that predictive targeting arguably comes with severe ethical and data protection issues that could lead to moral claims against the *advertiser*, it should be debated whether an opt-out option should also be offered to the advertisers.
- 5 Similar mechanisms are available for other platforms, such as, for example, *Google Adwords*, which offers ‘Customer match’, or *Twitter*, which automatically matches ‘custom audiences’ for its advertising customers according to their twitter handles (Twitter, 2022).
- 6 This reverse application of statistical inference, where knowledge about cohorts is projected back to information about individual cases, is at the heart of the ethically questionable epistemology of predictive reasoning (see Joque, 2022; Mühlhoff, 2021).
- 7 On the stronger provisions of the California Consumer Privacy Act (CCPA), see Blanke (2020).
- 8 Arigo et al. (2018), Oesterle et al. (2018) and Kamp et al. (2019) raise the issue that social media platforms might not cover some crucial criteria for research, for example, demographic attributes. Other biases are likely and need further scrutiny.

References

- Andreou A, Silva M, Benevenuto F, et al. (2019) Measuring the Facebook advertising ecosystem. In: NDSS 2019 - Proceedings of the Network and Distributed System Security Symposium, 2019, pp.1–15. DOI:10.14722/ndss.2019.23280.
- Andreou A, Venkatadri G, Goga O, et al. (2018) Investigating ad transparency mechanisms in social Media: A case study of Facebook’s explanations. In: NDSS 2018 - Network and Distributed System Security Symposium, 2018, pp.1–15. DOI:10.14722/ndss.2018.23204.
- Arigo D, Pagoto S, Carter-Harris L, et al. (2018) Using social media for health research: Methodological and ethical considerations for recruitment and intervention delivery. *Digital Health* 4: 1–15. DOI:10.1177/2055207618771757.
- Baldwin-Philippi J (2017) The myths of data-driven campaigning. *Political Communication* 34(4): 627–633.
- Barbu O (2014) Advertising, microtargeting and social media. *Procedia – Social and Behavioral Sciences* 163: 44–49.
- Barocas S and Selbst AD (2016) Big data’s disparate impact. *California Law Review* 104: 671.
- Beauchamp TL and Childress JF (2012) *Principles of Biomedical Ethics*, 7 Edition New York: Oxford University Press.
- Blanke JM (2020) Protection for “inferences drawn”: A comparison between the general data protection regulation and the california consumer privacy act. *Global Privacy Law Review* 1(2): 81–92.
- Bonevski B, Randell M, Paul C, et al. (2014) Reaching the hard-to-reach: A systematic review of strategies for improving health and medical research with socially disadvantaged groups. *BMC Medical Research Methodology* 14(1): 42.
- Borodovsky JT, Marsch LA and Budney AJ (2018) Studying cannabis use behaviors with facebook and web surveys: Methods and insights. *JMIR Public Health and Surveillance* 4(2): e48. DOI:10.2196/publichealth.9408.
- Caplan A and Friesen P (2017) Health disparities and clinical trial recruitment: Is there a duty to tweet? *PLOS Biology* 15(3): e2002040. DOI:10.1371/journal.pbio.2002040.
- CIOMS – Council for International Organizations of Medical Sciences and World Health Organization (2016) *International ethical guidelines for health-related research involving humans*. Geneva: CIOMS.
- Close S, Smaldone A, Fennoy I, et al. (2013) Using information technology and social networking for recruitment of research participants: Experience from an exploratory study of pediatric Klinefelter syndrome. *Journal of Medical Internet Research* 15(3): e48. DOI:10.2196/jmir.2286.
- Custers B, van der Hof S and Schermer B (2014) Privacy expectations of social media users: The role of informed consent in privacy policies. *Policy & Internet* 6(3): 268–295.
- Darmawan I, Bakker C, Brockman TA, et al. (2020) The role of social media in enhancing clinical trial recruitment: Scoping review. *Journal of Medical Internet Research* 22(10): e22810.
- Denecke K, Bamidis P, Bond C, et al. (2015) Ethical issues of social media usage in healthcare. *Yearbook of Medical Informatics* 10(1): 137–147.
- Ferguson AG (2016) Predictive prosecution. *Wake Forest L. Rev.* 51 HeinOnline: 705.
- Frandsen M, Thow M and Ferguson SG (2016) The effectiveness of social media (Facebook) compared with more traditional advertising methods for recruiting eligible participants to health research studies: A randomized, controlled clinical trial. *JMIR Research Protocols* 5(3): e161. DOI:10.2196/resprot.5747.
- Frandsen M, Walters J and Ferguson SG (2014) Exploring the viability of using online social media advertising as a recruitment method for smoking cessation clinical trials. *Nicotine & Tobacco Research: Official Journal of the Society for Research on Nicotine and Tobacco* 16(2): 247–251.
- Gelinas L, Pierce R, Winkler S, et al. (2017) Using social media as a research recruitment tool: Ethical issues and recommendations. *The American Journal of Bioethics* 17(3): 3–14.
- Guthrie KA, Caan B, Diem S, et al. (2019) Facebook advertising for recruitment of midlife women with bothersome vaginal symptoms: A pilot study. *Clinical Trials* 16(5): 476–480.
- Hildebrandt M and Gutwirth S (2008) *Profiling the European Citizen*. New York: Springer.
- Hitlin P, Rainie L and Olmstead K (2019) Facebook Algorithms and Personal Data. In: Pew Research Center. <https://www.pewresearch.org/2019/08/01/facebook-algorithms-and-personal-data/>

- pewresearch.org/internet/2019/01/16/facebook-algorithms-and-personal-data/.
- Hunter RF, Gough A, O’Kane N, et al. (2018) Ethical issues in social media research for public health. *American Journal of Public Health* 108(3): 343–348.
- InspireIM (2020) Online Patient Portal. <https://inspire-im.com/online-patient-portal/>.
- Joque J (2022) *Revolutionary mathematics: Artificial intelligence, statistics and the logic of capitalism*. London New York: Verso.
- Kamp K, Herbell K, Magginis WH, et al. (2019) Facebook recruitment and the protection of human subjects. *Western Journal of Nursing Research* 41(9): 1270–1281.
- King DB, O’Rourke N and DeLongis A (2014) Social media recruitment and online data collection: A beginner’s guide and best practices for accessing low-prevalence and hard-to-reach populations. *Canadian Psychology/Psychologie Canadienne* 55(4): 240–249.
- Kitchin R (2014) Big data, new epistemologies and paradigm shifts. *Big Data & Society* 1(1): 2053951714528481.
- Kosinski M, Stillwell D and Graepel T (2013) Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences of the USA* 110(15): 5802–5805.
- Kröger JL, Lutz OH-M and Ullrich S (2021) The myth of individual control: Mapping the limitations of privacy self-management. Available at SSRN. DOI:10.2139/ssrn.3881776.
- MD Connect (2017) Social Media & Clinical Trial Recruitment [Whitepaper]. https://cdn2.hubspot.net/hubfs/291282/documents/Gated_Content/White%20Paper%20-%20Clinical%20Trial%20-%20Social%20Media%20Patient%20Recruitment.pdf.
- Merchant RM, Asch DA, Crutchley P, et al. (2019) Evaluating the predictability of medical conditions from social media posts. *PLoS One* 14(6): e0215476. DOI:10.1371/journal.pone.0215476.
- Mühlhoff R (2020) Automatisierte Ungleichheit: Ethik der Künstlichen Intelligenz in der biopolitischen Wende des Digitalen Kapitalismus. *Deutsche Zeitschrift für Philosophie* 68(6): 867–890.
- Mühlhoff R (2021) Predictive privacy: Towards an applied ethics of data analytics. *Ethics and Information Technology* 23: 675–690. DOI:10.1007/s10676-021-09606-x.
- Mühlhoff R and Ruschemeier H (2022) Predictive analytics und DSGVO: Ethische und rechtliche Implikationen. In: Gräfe H-C (ed) *Telemedicus – Recht der Informationsgesellschaft: Tagungsband zur Sommerkonferenz 2022*. Frankfurt am Main: Deutscher Fachverlag, pp.38–67.
- Oesterle S, Epstein M, Haggerty KP, et al. (2018) Using Facebook to recruit parents to participate in a family program to prevent teen drug use. *Prevention Science* 19(4): 559–569.
- O’Neil C (2016) *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown.
- O’Neil C and Schutt R (2013) *Doing Data Science: Straight Talk from the Frontline*. Sebastopol, California, USA: O’Reilly Media Inc.
- Pasquale F (2016) *The Black Box Society*. Cambridge, Massachusetts, USA: Harvard University Press.
- Patel MX, Doku V and Tennakoon L (2003) Challenges in recruitment of research participants. *Advances in Psychiatric Treatment* 9(3): 229–238.
- Puffer S and Torgerson D (2003) Recruitment difficulties in randomised controlled trials. *Controlled Clinical Trials* 24(06): 214–215.
- Ryerson N (2021a) Advertising Clinical Trials on Facebook: Do’s and Don’ts. In: Antidote Blog. <https://www.antidote.me/blog/advertising-clinical-trials-on-facebook-dos-and-donts>.
- Ryerson N (2021b) How to Reach the Right Patients in Clinical Trial Advertising. In: Antidote Blog. <https://www.antidote.me/blog/how-to-reach-the-right-patients-in-clinical-trial-advertising>.
- Sedrak MS, Cohen RB, Merchant RM, et al. (2016) Cancer communication in the social media age. *JAMA Oncology* 2(6): 822–823.
- Semerádová T and Weinlich P (2019) Computer estimation of customer similarity with Facebook lookalikes: Advantages and disadvantages of hyper-targeting. *IEEE Access* 7: 153365–153377.
- Sitar S, Hartman B, Graham B, et al. (2009) P15-17. Social media as a tool for engaging and educating audiences around HIV vaccine research and clinical trial participation. *Retrovirology* 6(Suppl. 3): P218. DOI:10.1186/1742-4690-6-S3-P218.
- Sligh J, Abedtash H, Yang M, et al. (2016) A Novel Pipeline for Targeting Breast Cancer Patients on Twitter for Clinical Trial Recruitment.
- Statista (2022) Facebook: average revenue per user region 2021. <https://www.statista.com/statistics/251328/facebooks-average-revenue-per-user-by-region/>.
- The World Medical Association (WMA) (2013) Declaration of Helsinki – Ethical Principles for Medical Research Involving Human Subjects. <https://www.wma.net/policies-post/wma-declaration-of-helsinki-ethical-principles-for-medical-research-involving-human-subjects/>.
- Thornton RLJ, Glover CM, Cené CW, et al. (2016) Evaluating strategies for reducing health disparities by addressing the social determinants of health. *Health Affairs* 35(8): 1416–1423.
- Topolovec-Vranic J and Natarajan K (2016) The use of social media in recruitment for medical research studies: A scoping review. *Journal of Medical Internet Research* 18(11): e286. DOI:10.2196/jmir.5698.
- Tufekci Z (2015) Algorithmic harms beyond Facebook and Google: Emergent challenges of computational agency. *Colorado Technology Law Journal* 13: 203.
- Twitter (2022) Intro to Custom Audiences. <https://business.twitter.com/en/help/campaign-setup/campaign-targeting/custom-audiences.html>.
- van Dijck J (2009) Users like you? Theorizing agency in user-generated content. *Media, Culture & Society* 31(1): 41–58.
- Wachter S (2019) Data protection in the age of big data. *Nature Electronics* 2(1): –7.
- Wachter S and Mittelstadt B (2018) A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI. LawArXiv. DOI:10.31228/osf.io/mu2kf.
- Wertheimer A (2014) Non-completion and informed consent. *Journal of Medical Ethics* 40(2): 127–130. <https://www.jstor.org/stable/43282935>.
- Whitaker C, Stevelink S and Fear N (2017) The use of Facebook in recruiting participants for health research purposes: A systematic review. *Journal of Medical Internet Research* 19(8): e290. DOI:10.2196/jmir.7071.
- Wisk LE, Nelson EB, Magane KM, et al. (2019) Clinical trial recruitment and retention of college students with type 1 diabetes via social media: An implementation case study. *Journal of Diabetes Science and Technology* 13(3): 445–456.
- Wozney L, Turner K, Rose-Davis B, et al. (2019) Facebook ads to the rescue? Recruiting a hard to reach population into an Internet-based behavioral health intervention trial. *Internet Interventions* 17: 100246. DOI:10.1016/j.invent.2019.100246.