

Neuroscience, Mind Reading and Mental Privacy

Jesper Ryberg¹ 

Published online: 17 November 2016
© Springer Science+Business Media Dordrecht 2016

Abstract Many theorists have expressed the view that current or future applications of neurotechnology may prompt serious ethical problems in terms of privacy. This article concerns the question as to whether involuntary neurotechnological mind reading can plausibly be held to violate a person's moral right to mental privacy. It is argued that it is difficult to specify what a violation of a right to mental privacy amounts to in a way that is consistent with the fact that we usually regard natural mind reading as morally unproblematic.

Keywords Brain scans · Mental privacy · Mind reading neuroscience · Privacy rights

Introduction

Despite the fact that neuroethics has over the last few years attracted comprehensive attention and has as a field of research undergone an almost explosive growth, it is also a fact that this is an area of ethics which in several respects is still in its infancy. One of the implications of the youth of neuroethics is that it is easy to find proclamations of problems, or even threats, to which the development of modern neuroscience may give rise, which have not yet been followed up by a more thorough (sometimes not even preliminary) ethical scrutiny. I believe this is the case with regard to the implications which neuroimaging techniques may have in terms of privacy.

This is an original paper. No part of the paper has been previously published.

✉ Jesper Ryberg
ryberg@ruc.dk

¹ Department of Philosophy, Roskilde University, Post Box 260, 4000 Roskilde, Denmark

Several theorists have emphasized that current or future applications of neurotechnology prompt ethical challenges with regard to people's privacy. In fact, many have expressed the view that there are reasons for serious concern. For instance, Frank Tong and Michael Pratte have argued that 'mental privacy could face enormous new challenges, in both legal settings and beyond, as there has been no precedent for being able to look into the mind of another human being' (Tong and Pratte 2012, p. 502). Some theorists have even invoked traditional Orwellian parlance by contending that fMRI-based imaging technology may accomplish in real life what the Thought Police did in *Nineteen Eighty-Four*, or by holding that we are moving closer to 'Big Brother in your head' (see Boundy 2012; Gorman 2012). Irrespective of how precisely such views are presented, it is not difficult to understand the underlying drive of this concern.

We usually think of our minds as something to which only we ourselves have direct access or, as one commentator has put it, as the 'quintessential zone of privacy' (Shen 2013, p. 653). Thus, if parts of our mental life are suddenly made accessible to others by the use of neuroimaging techniques, then there seems to be reason to worry. However, it is also obvious that this more immediate reaction calls for clarification and further consideration. And on this point the academic literature is still sparse. The goal of this article is to contribute to the incipient discussion by critically examining one possible interpretation of why there could be ethical reasons for concern.

More precisely, the purpose in the following is not to consider how far the development of neuroimaging techniques have come with regard to mind reading except by noting the obvious fact that, given the nature of brain imaging techniques, the potential privacy problems related to this technology seem *practically* very different from the problems that arise from the use of other types of surveillance technology. After all, one does not just end up in an MRI-scanner without knowing this. And brain scanning usually requires a large degree of cooperation from the person who is being scanned (recall that even a slight move may produce noise in the pictures) (see e.g. Arstila and Scott 2011). However, this practical aspect, which I believe should itself put a damper on some of the more alarmist views on how neuroscience may invade our privacy, is not the issue here. Moreover, I shall not address the question as to whether there are some delimitations *in principle* with regard to what sort of access one person can possibly have to the mind of another (Gilead 2015). Rather, what I shall be concerned with is the purely ethical question as to whether neurotechnological mind reading can plausibly be held to violate a right to mental privacy. What I will argue is that it is very difficult to give content to what such a right consists in, and that, if the application of neuroimaging techniques constitutes a moral problem, or for that matter a problem in terms of privacy, this does not amount to a violation of this sort of right.

In order to reach this conclusion, the paper will proceed as follows. In section [Preliminary Conceptual Points](#), the path is cleared for the ensuing discussion by addressing a few conceptual and otherwise clarificatory points regarding the nature of a right to mental privacy. Subsequently, in section [Neurotechnological Versus Natural Mind Reading](#), attention is directed to the fact that in a certain sense we are all mind readers; it is argued that this fact confronts adherents of the view that

neuroscientific mind reading constitutes a problem in terms of mental privacy, with the challenge of identifying a morally relevant difference between neurotechnological mind reading and natural mind reading. The sections [Direct Versus Indirect Access](#), [Reliability](#), [Loss of Control](#) and [The Amount of Information](#) consider a number of possible ways in which this challenge could be met, arguing that, at the end of the day, none of the answers turn out to be successful. Finally, the [Conclusion](#) summarizes the main argument and adds a few remarks on the scope of the suggested conclusion.

Preliminary Conceptual Points

Whether it makes sense to hold that neurotechnological mind reading may infringe on a right to mental privacy is obviously not a question that allows for an immediate answer. On the contrary, in order to assess the idea of this sort of right violation and, not least, to avoid misunderstandings concerning the scope of the discussion, it is necessary to address a few preparatory considerations on what we are talking about in the first place.

The first thing that should briefly be mentioned concerns the nature of the right to mental privacy. First of all it should be underlined that what we are considering here is not the existence of a *legal* right. Over the last few years growing attention has been directed to the question as to whether there exists legal protection against mind reading with neurotechnology. In a US context, several legal scholars have asked if there is sufficient Fourth and Fifth Amendment protection of mental privacy. For instance, Farahany has recently argued that '[m]ental privacy is not sacrosanct under either the Fourth or Fifth Amendment, which provide procedural safeguard but not substantive ones to adequately protect mental privacy' (Farahany 2012, p. 408). Other commentators have held somewhat different views on the matter.¹ However, I shall not discuss here whether there is, or should be, a legal right to mental privacy. As is well known, legal rights can be morally justified in various ways and need not reflect underlying moral rights.² The purpose here is only to consider whether the use of neurotechnological mind reading may conflict with the existence of a *moral* right to mental privacy.

Furthermore, when considering the existence of a moral right it is well known that there are different views on what such a right consists in and, not least, what constitutes the relation between rights and duties. What I shall be assuming in the following is, roughly, that a moral right can be perceived as providing a protective zone around the person—that is, *in casu* the person's mental life—by prohibiting others from transgressing this zone. In other words, it is assumed that a right works as a moral constraint which means that a party who violates this constraint will be acting wrongly. Whether this constraint should be characterized in threshold or

¹ For a fine overview of different positions in the legal debate, see Shen (2013, pp. 694–97).

² For instance, recall Mill's famous contention that 'to have a right ... is ... to have something which society ought to defend me in the possession of. If the objector goes on to ask, why it ought? I can give him no other reason than general utility' (Mill 1962, p. 309).

absolutist terms is not important here. But it seems reasonable to believe that a justified violation of this right will at least require rather strong countervailing reasons.

The third thing that requires a preparatory comment concerns the kind of behaviour being assessed in terms of a privacy right. As the modern discussion of privacy has clearly revealed, there are many different ways in which acts may intrude on a person's privacy. However, what I shall be considering is limited to the question as to whether neurotechnological mind reading constitutes a privacy problem. The question of whether the fact that some party passes on information about a person's brain or mental life to a third party—a question which is highly urgent in many contexts such as in relation to the protection of health-care information—will not be considered here. As indicated, what I will consider is only whether there is a right to mental privacy such that someone who gains access to a person's mental life, that is, who engages in some sort of mind reading activity, against the will of this person, will be violating his/her right to privacy. Broader aspects of the privacy implications of the stream of information to other parties will not be addressed in this paper.³

The fourth point that should be underlined before embarking upon the discussion concerns another aspect of the scope of the moral assessment. Many theorists have underlined that we often have a strong interest in keeping parts of our mental life to ourselves. This factual claim is, I believe, indisputable and is certainly not something that will be questioned in the following.⁴ Moreover, and more importantly, I shall not here consider the broader question as to whether at times there is something wrong in engaging in activities that compromise a person's interest in concealing parts of her inner life. The wrongness of such activities may of course be explained in many different ways. For instance, it could be the case that other sorts of rights are being violated (e.g. a right to physical integrity) or that the wrongness is explained in terms of ethical theories that do not at all invoke the idea of moral rights (e.g. standard consequentialist theories). Whether there is something wrong in bringing about involuntary neurotechnological access to another person's mental life and, if so, when this is the case, are much broader questions which cannot be answered here. As indicated, the proceeding discussion is limited to the narrower question of whether the possible wrongness of such activities can be derived from a right to mental privacy; which is how the worry of the application of neuroimaging technology has often been framed.

Finally, it should be mentioned that the discussion is not bound to a narrowly defined view on the contents of our mental life. The discussion of neuroimaging as a method of mind reading usually comprises many different aspects of our inner life including, for instance, different kinds of preferences, biases, moods, thoughts, recognition of something, and mental activities such as lying. In the following, all these affective and cognitive elements of our mind will be regarded as parts of our mental life.

³ For an interesting discussion of the possible effects of mind transparency on our interpersonal transactions and relationships, see Richmond (2012).

⁴ For a general discussion of the significance of our privacy interests, see e.g. Moore (2003).

Neurotechnological Versus Natural Mind Reading

In the light of the previous conceptual points we can now turn to the main issue, namely, whether involuntary neurotechnological monitoring of a person's mental life can plausibly be held to violate a person's right to mental privacy. As mentioned, there is no reason to doubt that we often possess significant interest in concealing parts of our mental life from others. However, though there is an understandable tendency amongst theorists, who have directed attention to the privacy threats posed by modern neurotechnology, to emphasize the strength of this interest, it is also clear that this is not sufficient to justify the existence of a right to mental privacy. A plausible notion of a right to mental privacy presupposes that one engages in considerations on whether such a right can be articulated in a way that succeeds in providing the requisite sort of protection without becoming too inclusive. More precisely, a coherent account of such a right should not turn behaviour which we regard as morally unproblematic into right violations. By analogy, if an account of what it means to have a right to privacy in public implies that a person, who casts a simple glance at someone in the street in the casual and momentary way in which people usually watch each other in public areas, ends up as a right violator, then this account must surely be regarded as morally dubious. A plausible account of a right to privacy must provide protection without placing unproblematic ways of acting into a category of wrongful behaviour (see also Ryberg 2008).

But is this requirement of not being too inclusive relevant in relation to the discussion of a right to mental privacy? The answer is clearly in the affirmative. Even though we, as mentioned, often think of our mental life as something to which we have privileged access it is also a fact that others do, to some extent, have access to our minds. What is going on in our 'inner citadel' is *de facto* often not so private. One of the things that characterizes 'normal' people (i.e. people not suffering from cognitive or developmental impairment) is that they possess what is standardly referred to as a 'theory of mind'. That is, on the ground of observed behaviour (or sometimes lack of behaviour) we are able to make inferences about another person's unobservable mental states. In fact, we possess a remarkable ability to make successful judgements on others' beliefs, knowledge, intents, desires, moods, etc.⁵ In that sense we are all, as has been said, 'natural born mind readers' (Shen 2013, p. 656). But if that is the case, then it also becomes clear that the idea of a right to mental privacy may not be straightforward after all. In the same way as it would be absurd to hold that a person casually watching another person in the street is acting wrongly by violating this person's right to privacy in public, it seems equally absurd to suggest that we are all right violators when we in our daily life with other people—and independently of whether these people accept it or not—form judgements on what they think, intend, feel, or whatever else is going on in their minds. But if this is correct, that is, if we cannot plausibly be held to be right

⁵ Precisely how such judgements of others' minds are produced is a complicated question within philosophy of mind, and one that cannot be addressed in the present context. All that matters here is that people usually possess a theory of mind.

violators all the time we spend with other people, and if, at the same time, it is also believed that the use of neurotechnology as a method of getting access to other people's mental life may violate a right to mental privacy, then it has to be made clear what constitutes the difference between these two ways of approaching other people's mental life. In short, a plausible notion of a right to mental privacy must be specified in a way that succeeds in accounting for the morally relevant difference between what I shall in the following, for reasons of ease in exposition, refer to as respectively *neuroscientific* and *natural* mind reading. Unless this is possible, the idea of violations of a right to mental privacy will end up as too inclusive in the sense of comprising far too many acts in the scope of wrongful behaviour, and should—I suggest—consequently be rejected.⁶

In order to avoid misunderstandings, it should once again be underlined that the point here is not to consider the question as to whether there is anything wrong with involuntary neurotechnological mind reading. Thus, I shall not consider whether there are any possible moral differences at all between neurotechnological and natural mind reading. The purpose is only to address the narrower question as to whether there are any differences that can plausibly be used to explicate the content of the right to mental privacy. For instance, suppose *ex hypothesis* that the use of neurotechnological mind reading would generate a highly counterproductive aversion in people to modern technology. Obviously, such an aversion may constitute a reason against the use of this sort of mind reading, but it would hardly help in the specification of why neurotechnological mind reading should constitute a problem that has anything to do with privacy. In other words, given the fact that the parallel to natural mind reading clearly illustrates that it is not sufficient merely to contend that one violates another's right to mental privacy by reading his or her mind, some sort of more precise specification of when such a right is being violated will be required. Therefore, if it turns out that it is not possible to identify a morally relevant feature that can be used in the specification of why neurotechnological mind reading violates such a right (while natural mind reading does not), then what follows is not the general conclusion, that there is no moral problem in this type of mind reading, but only the more narrow conclusion that if there is a moral problem it does not consist in violations of a right to mental privacy.

But is it really a genuine challenge to specify a morally relevant difference which one might draw on in a more precise account of what a right to mental privacy consists in? Is it not pretty obvious that there are major differences between natural mind reading and neurotechnological mind reading, and that these differences can be adopted in the specification of when the right to mental privacy is being violated? Given that initially mentioned frequent subscription to the view that neurotechnological mind reading constitutes a serious problem in terms of privacy, I guess many people's immediate answer will be in the affirmative. However, as we shall now see,

⁶ In order to avoid misunderstandings, it should be mentioned that the point here is not to suggest that natural mind reading can never be carried out in a way that we would regard as wrong (or inappropriate), but rather that it is difficult to give content to the idea of a right to mental privacy in a way that succeeds in explaining the wrongness of neurotechnological mind reading without turning *unproblematic* instances of natural mind reading into activities which are wrong.

it turns out to be an intricate task to specify what such a morally relevant difference should consist in.

Direct Versus Indirect Access

Probably, one of the most immediate answers that come to mind is that the difference between neurotechnological and natural mind reading lies in the way in which the mental states of another person are being accessed. In the current neuroethical discussion of brain modifications, several theorists have underlined that there exists a crucial difference, for instance between, on the one hand, pharmacological ways of changing someone's mental life and, on the other, changes caused by addressing the person through arguments, instructions, or the like. And, that this is so even if at the end of the day the result—that is, the way in which the mental life of a person has been changed—is identical in both cases. The difference between the two approaches is usually characterized by the distinction between *direct* and *indirect* interventions. Now, though the issue here is not brain interventions but mind reading, could one not nevertheless trade on this distinction by contending that the difference between neurotechnological and natural mind reading lies in the fact that the former constitutes a more direct route to the mind of someone while the latter constitutes only an indirect approach to a person's mental life?

Though, as we shall return to shortly, there may be something to this distinction, it is also clear that in the absence of further qualifications the answer must be in the negative. It is obvious that natural mind reading constitutes an indirect approach to another person's mental life. The move of a hand, the accentuation of a certain word, or a certain facial expression, may be what lead us to infer what the person we are confronted with is thinking, intending, feeling, etc. However, it is equally clear that a neurotechnological approach to a person's mind cannot plausibly be regarded as direct. In so far as neuroimaging of the brain can tell us anything about what is going on in a person's mind, this is surely a route mediated by various technological and theoretical steps. In that sense it seems more appropriate to talk about another *indirect* method of mind reading. Be that as it may, it might still be relevant to examine whether the distinction between these two approaches should be regarded as morally relevant. In a recent defence of the moral significance of the direct/indirect distinction in relation to brain interventions, Bublitz and Merkel have tried to specify more precisely what this distinction amounts to. They suggest that '[d]irect interventions are those working directly on the brain (e.g. DBS, psychoactive substances) whereas ... indirect (or external) interventions are those stimuli which are perceived sensually ... and pass through the mind of the person, being processed by a host of psychological mechanisms', and, furthermore, that 'direct interventions can be primarily understood as electro-chemical or physical reactions following the laws of nature whereas indirect interventions involve psychological laws (or dynamics) ...' (Bublitz and Merkel 2014, p. 68). Again, we are not here considering brain interventions. But the question is whether the core idea could be maintained by holding that natural mind reading—e.g. inferences

from body language or facial expressions—is mediated through certain ‘psychological mechanisms’ which are clearly bypassed when the information is derived from direct observance of someone’s brain. And if so, could this distinction be invoked in the specification of the right to mental privacy and, thereby, explain why only neurotechnological but not natural mind reading should be regarded as morally unacceptable?

Whatever the outlined distinction between different ways of approaching other people’s mental life may have to offer in relation to brain interventions, I think it is clear that on closer scrutiny this has nothing to offer in the present context. Recall that the question we are addressing is whether there is a morally relevant difference between neurotechnological and natural mind reading that can explain why only the former approach constitutes a right violation. However, merely to hold that the one kind of access happens by watching brain processes while the other kind of access is given (roughly) through psychologically mediated physical manifestations of the mental life, cannot plausibly be held to identify the morally relevant difference. Rather, it seems that this suggestion is simply repeating the fact that initiates the moral question, namely, that we are considering two different ways of getting access to another person’s mental life. In other words, we are still confronted with the question of whether there is a morally relevant difference between the two factually different approaches to other people’s mental life which the adherent of a right to mental privacy might draw on.

Reliability

Even though the morally relevant difference between neurotechnological and natural mind reading cannot be explained merely by pointing to the fact that the two ways of accessing other people’s minds are different, it could well be the case that the morally relevant difference follows in a straightforward way from this factual difference between the two methods. It may be the case that the one method is simply more reliable than the other. Consider a simple (and admittedly somewhat rough) analogy. Suppose that we are considering the size of the monthly salary a person receives from her work. One approach might be to make the estimate on the ground of this person’s financial behaviour (consumption of clothes, food, etc.). Another possibility is to gain access to the person’s bank account statements. Though the latter method does not lead directly to the money this person earns, it is obviously preferable for the simple reason that it provides much more reliable information (in fact, presumably a correct account) of the person’s income. Could we, analogously, hold that the relevant difference between neurotechnological and natural mind reading can be boiled down to the fact that, while a theory of mind is basically a *theory* which may well (and in many cases does) leave us with incorrect judgements of another person’s mental states, the neurotechnological approach is characterized by the fact that it gets things right. In other words, could the fact that involuntary neurotechnological monitoring of someone’s mental life violates a person’s right to mental privacy, while this is not the case if one relies only on

natural mind reading, be explained in terms of the reliability of the judgements delivered by the former method?

If one believes that there are reasons to be worried about the possibility of neurotechnological monitoring of someone's mental life, then this could perhaps be explained by the belief that the direct approach (or, as we have just seen, what is more properly regarded as a different indirect approach) provided by neurotechnological methods will constitute an instance of genuine mind reading, that is, as something close to an infallible access to another person's mind. However, clearly it is not so simple. This attempt to clarify when mind reading constitutes a privacy right violation faces at least two challenges.

The first thing that should be noted is that, given the fact that there are no infallible methods of reading another person's mind, the idea of distinguishing between what is privacy violating and what is not privacy violating in terms of reliability will commit one to some sort of specification of when a mind reading method is sufficiently certain to constitute a privacy right violation. It is difficult to see how this challenge could be met in a way that does not seem highly arbitrary.

Secondly, and more importantly, it should be kept in mind that there is no simple relation between what type of mind reading we are considering and the degree of reliability of our judgements on other people's mental states. Some of the neurotechnological types of mind reading which have received comprehensive attention in the recent discussion, such as fMRI-based lie detection, are widely recognized as not yet being sufficiently effective to be applied in real-life cases (see e.g. Ryberg 2014). But even if such methods are significantly improved in the future it is very hard to believe that they will not still imply a non-ignorable error rate. On the other side, it is clear that, even though there is always the risk of mistakes when we make judgements about other people's mental lives, it is pretty clear that there are cases where we are quite good at it. If I have just seen someone eating a biscuit and in the next minute see and hear this person strongly denying this to be the case when another person asks her, then it is certainly not a wild guess that the biscuit-eater is lying. More generally, we are often—particularly in relation to someone we know well—fairly good at making judgements on what another person in a particular situation is thinking, feeling, or the like. What this implies is that if, as initially assumed, it is not the case that natural mind reading constitutes a violation of a right to mental privacy then it is hardly possible to explain in probabilistic terms why involuntary neurotechnological mind reading should constitute a privacy right violation.

Loss of Control

The distinction between direct and indirect access to another person's mind does not, as we have seen, succeed in providing a plausible account of why involuntary neurotechnological mind reading constitutes a problem in terms of mental privacy. However, even though the distinction in itself does not constitute a privacy-relevant moral difference, there may nevertheless be a point in maintaining the distinction. As mentioned, Bublitz and Merkel emphasize the significance of this distinction in

their discussion of neurointerventions, not because they believe that the distinction *per se* is morally relevant but because it is important in relation to what they regard as crucial, namely, the idea of mental self-determination. What they contend is that indirect interventions through stimuli which are sensually perceived make it possible for the person who is being influenced to react to these stimuli; or as they put it '[p]ersons have most control over interventions whose sensual substrates they perceive, particularly those rising to the level of conscious awareness' (Bublitz and Merkel 2014, p. 69). Though they concede that there may be grey areas, that is, that control may be reduced in cases where stimuli are being subconsciously processed (subliminal stimuli), the main point is that direct brain interventions imply a total loss of control. Now, even though Bublitz and Merkel are dealing with different types of neurointervention, it might be worth considering whether this idea could be extrapolated to the discussion of mind reading. In short, could the relevant difference between natural and neurotechnological mind reading be explained in terms of loss of control?

The idea that there exists a close relation between privacy and control is far from new. On the contrary, one of the most influential general theoretical approaches to privacy is *the control theory* which in one way or another specifies privacy in terms of control over access or information. For instance, Parker provides a version of a control account when he holds that 'privacy is control over when and by whom the various parts of us can be sensed by others' (Parker 1974, p. 281) and, in the same vein, Moore has more recently held that a 'right to privacy can be understood as a right to maintain a certain level of control over the inner spheres of personal information and access to one's body ...' (Moore 2003, p. 218). As should be clear, the purpose here is not to assess the plausibility of the control theoretical approach to privacy in general, but to keep the focus closely on mental privacy in relation to neurotechnological mind reading. Thus, the question is, how can the notion of control help in the identification of a morally relevant difference between natural and neuroscientific mind reading?

Drawing on Bublitz and Merkel's considerations, the answer could be that when it comes to natural mind reading we get access to another person's mental life by making inferences on the ground of the ways in which this person looks, reacts, behaves, etc. These physical manifestations of the person's inner life have in one way or another been processed by the person's psychological system and, therefore, are to some extent controllable by the person herself. In contrast, when the person's mind is approached directly (or more precisely, indirectly through the monitoring of the brain), the person is simply deprived of the possibility of controlling or influencing the stream of information that floats from her mind. As an illustration, consider again the parallel discussion about privacy in public. As part of his attempt to explain why a casual and momentary glance at someone in the street cannot constitute a violation of a right to privacy, while monitoring in a public area by the use of CCTV surveillance does constitute a violation of this right, von Hirsch has underlined that the latter kind of monitoring—which we may not know is taking place and which is conducted by unobservable observers whom we do not know anything about—does not make it possible to adjust our behaviour in the same way as we can when we are passing someone in the street or have a clear idea of who is

watching us at a certain time and location (von Hirsch 2000). In precisely the same way, it could be held that when it comes to natural mind reading it is, at least to some extent, possible for the person to control the physical manifestations of the parts of her mental life she wishes should remain unexposed to others. For instance, the person who is lying about having eaten the biscuit may try to look absolutely innocent; a person who is in a bad mood may try to camouflage this by smiling and in other ways appear happy; and the person who has a strong preference for something may do whatever she can to hide this from others. Conversely, when it comes to neurotechnological mind reading this way of getting access to a person's mental life seems to totally bypass the person's possibility of influencing the degree of accessibility to her inner life. Thus, have we here identified a morally relevant difference between natural and neurotechnological mind reading in terms of which it is possible to specify what it means to violate someone's right to mental privacy?

There is no doubt that this suggestion seems more appealing than the answers considered in the previous sections. And, as we have seen, by placing the focus on a person's possibility of controlling access or information, the answer accords with well-established general approaches to privacy. However, on closer scrutiny it turns out that this approach does not succeed in providing a ground for the moral distinction between neurotechnological and natural mind reading.

Consider first neurotechnological mind reading. How does this way of getting access to the mind of another person deprive this person of the possibility of adjusting to the fact that she is being monitored or of controlling what the mind reader gets access to? The answer seems to depend upon what part of the person's mind is being monitored and on how precisely this is done. Suppose first that a person against her will is placed in a fMRI scanner in order to get an insight into what she is thinking. To hold that the person in this case will not be able to influence the result of the mind reading process is clearly false. As anyone who is familiar with this type of imaging technology will know, it is very easy to obstruct brain scans. As mentioned, even a slight move, such as moving of the tongue, may produce a devastating noise at the pictures. Moreover, suppose that we ignore this simple way of influencing the result of a brain scanning, for instance, by imagining that fMRI-technology was improved in a way that made it insensitive to movements or, even more hypothetically, that it was possible to monitor the mental life of a person thorough an implanted brain device or perhaps by using some sort of remote brain scanning system. Now, even under these conditions it is not clear that the individual whose mind is being monitored is deprived of the possibility of influencing the process or the result. For instance, if we imagine that it is the *thoughts* of the person that are being monitored, then it certainly seems that the individual may well have some influence on what is being monitored. The person can to some extent make decisions on the content of her thoughts. There may be other parts of our mental life which we have much less control over, but it is remarkable if the account of privacy implies that involuntary reading of a person's thoughts does not constitute a privacy violation; after all, many seem to regard thought reading as the most extreme example of privacy invasion. To this it might perhaps be replied that the possibility of influencing an instance of thought reading is at least removed in cases where monitoring of the brain occurs in secret.

However, I shall not here enter into a detailed discussion of this possibility. Sufficient has been said to indicate that what it means to specify violations of a right to mental privacy in terms of deprivation of control is not straightforward. Moreover, the most problematic aspect of this account, as we shall now see, relates to the implications associated with natural mind reading.

How does the suggested interpretation of violations of the right to mental privacy accord with our initial assumption that natural mind reading does not constitute a right violation? As indicated, there is no doubt that there are many cases where it makes sense to hold that we possess control over the access others have to our inner life. However, this is not the whole story. First, there are several cases where it is not possible for us to control the physical manifestations which make it possible for others to infer what is going on in our minds. It may well be impossible for someone in a particular context to avoid certain facial reactions to what is happening, such as, for instance, blushing with embarrassment, trembling with excitement or rage, shaking from nervousness, or numerous other bodily expressions of emotions. Second, the inferences we draw on others' mental lives are often very tightly bound up with the context of the situation. As mentioned, if I have just seen a person eating a biscuit, who denies this the next minute, there is very little that person can do to prevent me from drawing the conclusion that she is now engaged in the mental operation (whatever this precisely consists in) of lying. Third, in cases where we are watching a person (for example, a family member or a bus passenger) who is not aware of this; this person's possibility of controlling the inferences I make on what is going on in her mind is obviously very limited. What these examples indicate—and I believe many more could be added along the same lines—is that, if we maintain the view that natural mind reading does not violate a right to mental privacy, then the idea of explaining violations of such a right in terms of control cannot be regarded as satisfactory.

Thus, where do the previous considerations lead us? It seems clear that the notion of a right to mental privacy, which is based on the idea of control of access, does not provide a plausible account of the morally relevant difference between neurotechnological and natural mind reading. There are cases where we can control access to our mind, even when our mind is approached with neurotechnological tools and, perhaps more importantly, there are several cases of natural mind reading that leave an individual with little control. Moreover, it seems that, once one seeks to explain more precisely what sort or degree of control it is that matters with regard to violations of mental privacy, one may easily end up in the unfortunate position where an account of control, designed to explain why all or most instances of neurotechnological mind reading are wrong, will the more easily result in including unproblematic cases of natural mind reading in the scope of rights violations. While, oppositely, an account designed to keep natural mind reading morally unproblematic will leave it difficult to capture the instances of neurotechnological mind reading which seem problematic. Thus, in the present form the suggested account does not appear to place us on the right track.

The Amount of Information

As we have seen, the main problem confronting attempts to explain how neurotechnological mind reading can plausibly be said to violate a right to mental privacy consists in the fact that the different accounts become too inclusive, in the sense that they result in turning morally unproblematic instances of natural mind reading into right violations. For instance, an exposition of the right to mental privacy that implies that it is wrong to read someone's thoughts by neurotechnological means will face the challenge that it is sometimes possible for us to make judgements on others' thoughts without the use of such means. However, perhaps there is nevertheless an important difference. Even though natural mind reading can sometimes provide us with an insight into what is going on in someone else's mind, it is an indisputable fact that there are many cases where this is not possible. Thus, when the problem arises of identifying the morally relevant difference between neurotechnological and natural mind reading, does the key to the answer not lie in the fact that neurotechnological mind reading can provide us with much more information than we get through natural mind reading? In other words, following this idea, the reason that neurotechnological mind reading constitutes a problem in terms of mental privacy is not that this technology has slightly opened a usually closed door to the mind of others, but rather that the door has been pushed wide open.

At first glance the idea that neurotechnological mind reading does not consist in something with which we are totally unfamiliar but rather lies in the *capacity* related to this technology, is appealing. Moreover, this idea has an affinity with other sorts of privacy considerations. For instance, in so far as CCTV surveillance in public areas constitutes a privacy problem, it might be felt that this cannot be explained merely in terms of the fact that the technology makes it possible to watch people in the streets—after all, this is something we all do when we walk about in a public area—but rather in the capacity to watch people very closely and/or for long periods of time.

However, on closer inspection the idea of specifying why neurotechnological mind reading violates a right to mental privacy in terms of the capacity of this technology seems dubious. The obvious challenge is to specify what this increased capacity consists in. Leaving aside considerations of the reliability of neurotechnological mind reading relative to natural mind reading, the only two capacity-based ways of explaining why neurotechnological mind reading violates the right to privacy are in terms of either the type or the amount of information that is provided. Of these possibilities the first can easily be excluded. As mentioned, natural mind reading can provide us with information both of intentions, thoughts, moods, emotions, and other parts of our inner life. It is certainly hard to think of any type of information of our mind which is always inaccessible through natural mind reading but which can be obtained by neurotechnological means (and, even if there is such information, it is probably dubious that adherents of a right to mental privacy would be satisfied with limiting the scope of right violations to this particular type of information). But then, what of the second possibility? Is it more plausible to hold

that infringement of the right to mental privacy relates to the *amount* of information that is gathered?

The problem which, in my view, gives reasons to doubt this answer is twofold. First, it is necessary to say something about how degrees of information should be measured. This in itself is an intricate task (e.g. should reading of an individual's thoughts be measured in time, and what about the monitoring of mood or emotions?). Second, one would have to specify when the amount of information is sufficiently large to constitute a right violation (see also Ryberg 2007). Indeed, it is hard to imagine that this can be achieved in a way that will not leave the suggestion extremely arbitrary (e.g. to hold that you are not violating the right to mental privacy if you are mind reading another person for two minutes, but that two and a half minutes would constitute a right violation seems very arbitrary). Furthermore, independently of how precisely the two challenges are met, the overall theory one arrives at will still turn instances of natural mind reading, which we have initially assumed morally acceptable, into cases of wrongful behaviour. All together, these problems clearly indicate that, despite any immediate appeal, this approach does not succeed in delivering the sought-for answer.

Conclusion

As we have initially seen, it is not difficult to find proclamations in current academic literature (and indeed in non-academic writings) of the view that modern neuroscience prompts serious problems in terms of privacy. In this paper, I have tried to contribute to the discussion of this view by examining the contention that involuntary neurotechnological mind reading constitutes a violation of the right to mental privacy. The point of departure of the discussion has been, I believe, the uncontroversial assumption that, in order to sustain this contention, it is not sufficient merely to show that the view nicely explains the wrongness of cases of neurotechnological mind reading that strike us as morally unacceptable. It is also necessary to consider whether this right is consistent with the fact that natural mind reading—that is, our everyday inferential judgements on other people's mental lives—does not seem morally problematic, even when this is done against the will of those whose minds we are considering. What I have argued is that, on closer inspection, it is difficult to give content to the idea of a right to mental privacy in a way that succeeds in explaining the wrongness of instances of neurotechnological mind reading without turning unproblematic instances of natural mind reading into an activity that should be regarded as morally wrong. However, in order to avoid misunderstandings, it is important to emphasize what these considerations have not established.

First of all, it should again be underlined that I have not argued that there is nothing wrong in subjecting someone to involuntary neurotechnological mind reading. Whether there is something wrong in this sort of practice is a much broader question than the one addressed here, which has analysed the question only in terms of the right to mental privacy. Furthermore, even within this narrower focus it has not been shown that neurotechnological mind reading cannot plausibly be held to

violate a right to mental privacy. What I have done is to consider different accounts of what it means to violate the right to mental privacy and to show that, given the assumption that natural mind reading is not morally unacceptable, none of these accounts manage to provide a plausible explanation of the wrongness of neurotechnological mind reading. However, I cannot exclude the possibility that other explications of the right to mental privacy can be given. If the reader has a proposal: please publish!

References

- Arstila, V., and F. Scott. 2011. Brain reading and mental privacy. *Trames* 15: 204–212.
- Boundy, M. 2012. The government can read your mind: Can the constitution stop it? *Hastings Law Journal* 63: 1627–1644.
- Blublitz, J.C., and R. Merkel. 2014. Crimes against minds: On mental manipulations, harms and a human right to mental self-determination. *Criminal Law and Philosophy* 8: 51–77.
- Farahany, N.A. 2012. Incriminating thoughts. *Stanford Law Review* 64: 2012.
- Gilead, A. 2015. Can brain imaging breach our mental privacy? *Review of Philosophy and Psychology* 6: 275–291.
- Gorman, C. 2012. The mind-reading machine: Veritas scientific is developing an EEG helmet that may invade the privacy of the mind. *IEEE Spectrum*.
- Mill, J.S. 1962. *Utilitarianism*. London: Fontana.
- Moore, A.D. 2003. Privacy: Its meaning and value. *American Philosophical Quarterly* 40: 215–226.
- Parker, W.A. 1974. A definition of privacy. *Rutgers Law Review* 27: 275–296.
- Ryberg, J. 2007. Privacy rights, crime prevention, CCTV, and the life of Mrs Aremac. *Res Publica* 13: 127–143.
- Ryberg, J. 2008. Moral rights and the problem of privacy in the public: A reply to lever and goold. *Res Publica* 14.
- Richmond, S. 2012. Brain imaging and the transparency scenario. In *I know what you're thinking*, ed. S. Richmond, et al. Oxford: Oxford University Press.
- Ryberg, J. 2014. When should neuroimaging be applied in the criminal court? On ideal comparison and the shortcomings of retributivism. *Journal of Ethics* 18: 81–100.
- Shen, F. 2013. Neuroscience, mental privacy, and the law. *Harvard Journal of Law and Public Policy* 36: 653–713.
- Tong, F., and M.S. Pratte. 2012. Decoding patterns of human brain activity. *Annual Review of Psychology* 63: 483–509.
- von Hirsch, A. 2000. The ethics of public television surveillance. In *Ethical and social perspectives on situational crime prevention*, ed. A. von Hirsch, D. Garland, and A. Wakefield, 59–76. Oxford: Hart Publishing.