“I spy, with my little sensor”

Fair data handling practices for robots between privacy, copyright and security

"I Spy, with my Little Sensor": Fair Data Handling Practices for Robots Between Privacy, Copyright and Security

Burkhard Schafer[a], Lilian Edwards[b]

a) School of Law, The University of Edinburgh; Old College, Edinburgh, EH8 9YL; 0131-6502035; b.schafer@ed.ac.uk corresponding author
b) School of Law, Strathclyde University, Glasgow; lilian.edwards@strath.ac.uk

Abstract: The paper suggests an amendment to Principle 4 of ethical robot design, and a demand for "transparency by design". It argues that while misleading vulnerable users as to the nature of a robot is a serious ethical issue, other forms of intentionally deceptive or unintentionally misleading aspects of robotic design pose challenges that are on the one hand more universal and harmful in their application, on the other more difficult to address consistently through design choices. The focus will be on transparent design regarding the sensory capacities of robots. Intuitive, low-tech but highly efficient privacy preserving behaviour is regularly dependent on an accurate understanding of surveillance risks. Design choices that hide, camouflage or misrepresent these capacities can undermine these strategies. However, formulating an ethical principle of "sensor transparency" is not straightforward, as openness can also lead to greater vulnerability and with that security risks. We argue that the discussion on sensor transparency needs to be embedded in a broader discussion of "fair data handling principles" for robots that involve issues of privacy, but also intellectual property rights such as copyright. To balance respect for these rights with a need for security requires a framework that goes beyond the duties of the roboticists to an analysis of the duties of the public when interacting with a robot.

Keyword: design ethics; robot ethics; privacy; copyright, principles of robotics

Acknowledgement: Work on this paper was supported by CREATE, the RCUK Centre for Copyright and New Business Models in the Creative Economy

1.Introduction

This paper builds on Principle 4 of the the EPSRC Principles of Robotics[1]:

Robots are manufactured artefacts. They should not be designed in a deceptive way to exploit vulnerable users; instead their machine nature should be transparent.

We will argue however that it also need to be understood within the more general Principles 2 and 3:

2) Robots should be designed; operated as far as is practicable to comply with existing laws & fundamental rights & freedoms, including privacy.

3) Robots are products. They should be designed using processes which assure their safety and security.

---

1 https://www.epsrc.ac.uk/research/ourportfolio/themes/engineering/activities/principlesofrobotics/

We will argue that Principle 4 ought to be amended to include also other forms of deceptive or misleading design choices that can have an impact on rights of users or the general public. In particular, we argue that sometimes, misleading design choices can have a negative impact on privacy and other "data control rights" such as intellectual property rights. However, as we will see, some of these design choices are capable of furthering Principle 3, by increasing the chances of safe operation of robots. Balancing conflicting policy goals is a frequent problem in law, and we find in doctrinal law several tools and procedures for this task. This makes Principle 2 the bridging principle between the issues raised by Principles 3 and 4 – it requires us to find ways to operationalize at the design stage appropriate consideration of conflicting legal and ethical demands, privacy and security amongst them. Seen from this perspective, Principle 4 can be seen as just one aspect of a more general problem of "transparent" robotic design which does not just affect the outward appearance, but also the internal algorithms and data processing functions. The issue of algorithmic transparency and none-deceptive robotic design are interconnected, and not just a problem of deception about their machine nature, but a more generic question of how to balance transparency versus security in the design process.

One aspect of this balancing process is to expand the scope of the legal and ethical debate beyond the duties of the designer, and to ask about corresponding duties by owners, users and, indeed, third parties towards robots. The Principles intentionally avoided to include rights of third parties, also to avoid any implication that robots may be considered right holders. However, we do not mean here to enter the discussion regarding robots as potential holders of rights. Rather, the deliberative process that is needed to turn ethical principles into concrete design choices for pro-ethical and law compliant robot design has to include an acknowledgement of standards of "appropriate behavior around robotic devices" by people other than the designers. To decide whether a design choice is "safe" for instance inevitably involves an assessment of the behavior of people that use or otherwise interact with a robotic device. This then also means that we can't any longer determine the duties of a robot designer insolation from the standards of behavior that we can expect from other people who will interact with the machine. If, hypothetically, a designer could rely on other people not to break the law when interacting with robots (e.g., not to damage them), s/he could rely on legal deterrents and take design choices that promote values other than (physical) safety, such as for instance privacy. A quick example can illustrate the point:

Assume a self-driving robot relies for its safe operation on a number of sensors that tell it if it is approaching human beings. If its sensors indicate that a human could be hit, it decelerates. If vandals were to interfere with the sensors, e.g. by covering them in paint, this increases the risk of an accident. A design choice reducing this risk is to hide the sensors from plain view. This, as we will discuss in more detail below, can increase the privacy risks posed by the robot, as it may also prevent third parties from taking legitimate and risk-free countermeasures against being filmed and recorded by the machine. This creates a potential conflict between Principle 2 (privacy) and Principle 3 (safety). An ethically acceptable solution will have to ask if it is appropriate to rely on people to act in a lawful way towards the robot, and not to interfere with or damage the sensors. The burden of the designer to build ethically and legally sound machines can in part be shifted to a burden of third parties to interact with robots in particular ways. This "responsibility shifting" in turn raises both factual

and ethical questions. Demanding third parties to respect the physical integrity of a robot is ethically more likely to be sound, and factually more likely to be successful, if the robot in turn acts fairly towards these parties. In our scenario, we hypothesize that bystanders will be less likely to interfere with the sensors of a robot if they can be reassured that any information the robot has to gather about them in order to interact with them safely will be treated "fairly" (also in the sense of Data Protection Law), that is, kept to the minimum necessary, only used for the purpose of safe navigation, and destroyed once it is not any longer needed.

As this simple example shows, in actual practice it can be difficult to define the ethical and legal duties of a robot developer, in isolation from the duties and standards of behavior of third parties, the way the Principles do. Rather, we need more complex and interactive models, in the case at hand in particular a theory of "fair data handling practices" that allows a robot to record, copy and analyze all the data it needs for safe operation, while at the same time optimizes the design to protect the legitimate rights of the people whose data it needs. In the next section, we will discuss why sensors, and deception regarding the sensory capacities of robots, pose significant challenges that can only be addressed within such a wider theoretical approach.

2. Deception, Sensors and Privacy

Robots pose some unique challenges for fair data handling practices, challenges that are at least in part caused by their capacity to "deceive", if inadvertently, the people they interact with. For the purpose of this paper, we will not distinguish between misleading and deceitful behavior, and cover both forms of misdirection equally. Long before modern technology, humans developed privacy preserving techniques, from the curtains to the windows to the growing hedges, from learning when to whisper to washing away one's scent in a stream when hunting. These strategies protected them from the prying eyes of fellow man as much as from the interest of non-human predators. Furthermore, they protected not just privacy, but also other informational interests, including valuable information monopolies such as trade secrets or know-how. Being the only one to know how to make fire in a tribe, or to know where the best fruit was growing, was arguably a more important informational interest in early societies than privacy. Crucially, these not only protect information, but also the sharing and exchange of information – whispering is a way to protect your data while sharing information, sound proofing your studio protects your privacy and your commercial interests in your music, but by suppressing noise also allows you to record it in ways that can be more easily shared.

The walls we build around us do not just keep the warmth in and the rain out, but also information in and observers out. The law, with its system of rules and exceptions, frequently gave formal recognition to these low-technology protection measures. Hannah Arendt's (1958) distinction between the private, the public and the social tracks in many ways these physical architectures. Similarly, our law recognizes these efforts to create private spaces, with the house, the locked cupboard or the hedge-protected garden the archetype of "reasonable expectation of privacy" and "security in our houses, papers and effect". Building these types of barriers keep other people and the state out, and created very early on in our legal history the tort of "intrusion into seclusion" for those who ignore these physical and symbolic walls. (Kang 1998, esp p. 1202). But law does not just normatively reinforce physical boundaries that

keep others out, it also recognizes our efforts to create private spaces as means to more freely collect and exchange data, to enable our choices in how we express ourselves to others and in this way form an identity (DeCew1997, esp. 77). From a legal perspective, the domestic purpose exemption in Article 3 of the European Data Protection Directive is a particularly interesting and relevant example in this respect. It permits individuals to collect the private data of others, provided this takes place within clearly identifiable spaces, i.e. the data collector's home, and for purely domestic or family purposes. An example could be to take photographs of guests for a family album, or give an au pair a rota of visitors and the schedule of the children's whereabouts. With robots as new domestics however, the level of intrusiveness, the persistence of data collected and the ability to share it potentially increases dramatically.

Generally, robotics threatens to render these low-tech solutions to the privacy problem more and more redundant. We face an increasing range of sensor capacities, many of which we did not encounter, or did nor encounter in a significantly threatening way, in our evolutionary past (see e.g. Fradella et al 2010 p. 303 ff). I know that a wall protects my privacy also because I know that my neighbor can only detect information in the normal visual band, so erecting a wall or simply moving outside his line of sight is sufficient. If I can't know any longer if my robotic neighbor also senses heat, or smell, this becomes inefficient. Evolutionary acquired and habituated privacy preserving strategies can thus easily be undermined.

Robots also are increasingly mobile and ubiquitous. For many applications, we will (have to) "invite them into" our home. This includes medical care robots, but also entertainment and service devices. Just as the Victorian upper class took it for granted to be served and surrounded by an army of servants, we may face a situation again where rather than operating "dumb" machinery, we will re-invite armies of robo-servants into our homes, and just as with their Victorian counterparts, we will need to find appropriate "rules of engagement" with them (Hamill 2006 p. 245ff). As the Victorians knew though, nobody is a hero to his domestics. Inevitably, information is disclosed to them that makes the subject of the information very vulnerable. But at least, with domestics the lord or lady of the house could anticipate what exactly they would able to see, they would understand the normative (both social and legal) environment that restrained them from collecting and most importantly sharing data about their employer. The understanding of the normative environment together with the understanding of the sensory capacities enabled rational risk assessment and management – you could probably trust your butler with your dirty underwear, but maybe not with a blood-drenched shirt. In the bedroom, one needn't worry about the heat signature when entertaining a companion, but one would possibly chose to keep the noise down. All this while relying on the butler to knock first before entering the bedroom, making the closed door both a physical but also a symbolic privacy protecting device.

Robotics threatens these defensive strategies not just because they can use sensors outside the visual or acoustic spectrum of humans, or because of their mobility that allows sensing in spaces previously protected. In addition, when they imitate the outward appearance of humans, or indeed non-human animals, even in cases where their robotic nature is plain visible (as per principle 4), it can misdirect our efforts to shield ourselves form them. Admittedly, this claim would benefit from stronger

empirical backing, but the importance of "eye contact" with a humanoid robot has long been recognized as important for interaction and joint attention (see e.g. Yonezawa et al 2007). This indicates that external features that look like eyes are interpreted both as the space where a robot's sensors are placed, and also are used to ground mini-theories about their capacity. In short, if a robot seems to have human eyes, we interpolate form this its line of sight, its scope of vision and also tend to assume it senses within the normal visual spectrum. (see e.g. MacDorman et al 2005; Xu et al 2016). The Internet is abundant with clips of people "sneaking up" on Asimo or similar loosely anthropomorphic robots from behind – and while their sensors "may" indeed be located in their eyes, and have vision restrictions similar to a human, this may well of course be false.

Part of ethical design therefore should also be to indicate the sensory capacities of robots in ways that facilitate the emergence of "intuitive" defenses of the type we use with other humans, and refrain, where possible, from inviting misleading inferences. The ease which with we can avail ourselves of effective low-tech, low-cost defensive mechanism (such as "moving out of the line of sight") should also be a factor in the evaluation of intrusiveness, when a choice between different sensors can be made.
For the UK, this approach is also in line with the Information Commissioner's practice guide on CCTV, in particular body worn video cameras (BWV) which due to their mobility, multi-sensor capability and small size are the closest equivalent to robotic sensors:

"As BWV cameras can be quite small or discreet, and could be recording in fast moving or chaotic situations, individuals may not be aware that they are being recorded. It is therefore important that clear signage is displayed, for example on an individual's uniform, to show that recording is taking place and whether the recording includes audio"[2]

Data protection law is one driver behind this suggestion, but "fair sensing" practices go beyond personal data, let alone sensitive personal data. Humans protect not just data about themselves from the eyes of others, but also their business ideas, scientific or technological discoveries, or personal skills. Here too we reason instinctively about sensory capacities by potential adversaries. Even school children sometimes build a wall of books around them during exams, to prevent others from cheating and gaining an unfair advantage. A low tech solution to the intrusion of informational spaces, which companies or professionals replicate, in more sophisticated form, when their protect trade secrets or industrial IP in a locked safe.

Intellectual property law is therefore another legal constraint that needs to be observed under this header of the Principles, and a broader notion of "fair data handing practices" that goes beyond DP law may be needed. Robots can harm humans through their sensors not only if they collect personal data about them, but also when they make copies that document a human's knowledge, skills or information. As a straightforward example, a personal assistance robot hired to talk with its client about their day at work must not relay this sensitive trading data to its owner. Maybe a more futuristic example that moves us beyond (current) law into the

[2] https://ico.org.uk/media/for-organisations/documents/1542/cctv-code-of-practice.pdf

field of aspirational ethics, if an industrial robot observes my movements to improve collaboration and avoid collusions, and in the process learns enough to make my job less secure (the way I move is my unique selling point), this may violate a legitimate "interest" of mine that deserves protection or compensation.

In all of these examples, not misleading the user/owner about the sensory capacities is a minimal requirement to allow them taking appropriate defensive mechanisms. We need to (re)learn what we can safely talk about in our home, when the new "domestics" are surrounding us.

2. Beyond fair sensing duties

So far, we have suggested only a modest amendment to the Principles. Just as robots should not deceive vulnerable users as to their nature as machines, so they should not mislead users as to their sensory capacities. Or in short, do not unnecessarily give your robot protrusions that look like mammalian eyes looking forward, when in reality it senses 360° in the infrared spectrum. Alternatively, indicate in other ways what the real sensory capacities are, in a way that makes it easy for a person interacting with the machine to respond appropriately. The examples that we discussed so far raise however also questions that lead in a more radical way beyond the Principles. The Principles try to establish duties that the developers of robots owe to people who interact with their machines. Questions such as whether robots should under some conditions be given a capacity for self-defense, which would have established a more symmetrical relation of rights and duties, were brought up but ultimately not pursued at the time the original Principles were discussed. Nonetheless, the discussion above leads to a necessary nexus between duties that the developers owe, and possible legal or ethical duties that are owed to them or the robot owner.

As a simple example, to allow safe robot design may involve a duty for third parties to disclose or share certain information with the robot that in the past they could legitimately keep for themselves, or that was legally protected in case it was copied without permission. A care robot that navigates e.g. an art exhibition to assist its visually impaired owner may have to create images of the exhibits and share them with other robots present, simply to avoid running into the wall or the other machines. This can potentially create legal conflicts, if, for instance, gallery owner prohibits the taken of photographs (exercising his property rights), or if he evokes copyright law to prevent even this "incidental" copying. In some jurisdictions, copying for functional rather than expressive purposes is already permitted and so should not be a breach of copyright, something that enables copy-reliant technologies from reaching the market (Sag 2009). However, once several machines in the gallery co-ordinate their actions and in the process also *share* this data between them (to prevent e.g. them all moving towards the same image at the same time) even this line of argument may reach its limits. In other words, we face a potential conflict between the duty to build safe robots, and the duty to observe "information control laws" such as data protection or IP law. If third parties chose to not share certain information, safety may suffer as a result. Resolving this tension requires to talk not just about the rights of these third parties, but needs also a discussion if at the very least an ethical duty to "share with a robot the information it needs for safe operation" can be established.

In the above examples, the "owner" of the information remained passive. Citizens may however also chose to use technology to actively prevent sensors from noticing them (e.g. camouflage face paint - https://cvdazzle.com)/ As in the above discussion, this may mean that they incur a greater risk that the robot runs into them, as it may hinder its ability to identify humans in its environment. If we add machine learning and third parties into this setting, this can create even more complex legal issues. If my face provides a data point from which I know the robot learns, do I have "quality assurance duties", a duty to be a good example - in the same way one can argue that I have an additional ethical duty not to cross a street when the lights are red if I know I'm observed by a child that is still learning safe conduct. If as a third party, not contractually obligated or otherwise involved with the development of a robot, intentionally manipulate its learning to induce dangerous behavior, does this give rise to liability or even moves into the territory of the Computer Misuse Act? And finally, if I contribute to the learning of a machine, do I have a claim on what it produces as an outcome?

Traditional negligence law, its conceptualisation of the duty owed to "ones neighbours", the distinction between act and omission, and the concept of reliance liability will all be part of the legal answer *after* an accident happened. For the purpose of the discussion here however, the question is posed slightly differently: At the point of developing a robot, can/should the designers, in discharging their duty to build safe *and* law compliant machines:

1. Rely on the *ethical/social* duty by third parties not to manipulate the sensing or knowledge acquisition of the machine
2. Rely only on a narrower *legal* obligation to refrain from certain foreseeably dangerous manipulation of the robot and its sensors
3. Not rely at all on a cooperative environment when thinking about the safety and law compliance of the robot they build – after all, not all laws are observed by everybody.

To make clear why this issue arises in the context of a discussion on "sensor transparency": IF we accept the ethical obligation discussed above, i.e. that robots should normally disclose how and with that what they can sense, then they inevitably open themselves up to manipulation or interference. If we in addition accept 1, or at the very least 2, this is less of an issue than if we accept 3.

If we accept 2, then we have to deal head on with the issue of duties owed by third parties when interacting with robotic devices. On the one hand, it is hard to owe duties to non-humans, robots are not suitable examples of right holders, and nobody owes them a duty. Robots are, as the Principle states, mere artefacts. On the other hand, humans can of course encroach on the rights of *other humans* through the way they act towards objects owned by them. I owe a duty to my neighbour not to burn down his barn. The type of "information disclosure duties" that this paper discusses sit uneasily between these two clear issues. Do we owe a (legal) duty to an absent robot operator (or designer) not to confuse their robot's sensing and/or training? Under which conditions is it reasonably foreseeable a robot would be confused when we interact with its sensors? In the UK, the law does not recognise a duty not to lie to strangers when, for instance, giving them directions. Normally, such a duty is only triggered when there is a specific role that people have in virtue of their profession

(e.g. a medical or legal advisor) or a special relationship such as that of a parent to their child. But then again, sending a lost child asking for direction astray would be a different proposition, especially in countries that recognise a general duty to rescue. Are machines that "still learn" analogous to such a situation?

To recap: Robots potentially increase the threat to our privacy considerably. Their mobility, their potentially small size, and the range of sensors that they carry render some regulatory protections that worked, to a degree, with fixed CCTV cameras moot. Like Body Worn Video, they can follow individuals around and circumvent some protective measures by using an array of sensors. Unlike BWV however, they are only limited in their endurance by their battery live, and can thus also extend the temporal scope of surveillance. Furthermore, to reap some of the benefits that they offer, we need to "invite them into" our private spaces, e.g. as domestic or care robots. The law, with its emphasis on consent, loses much of its protective power under these conditions. In this environment, it becomes particularly important the we re-learn intuitive, privacy protecting behaviour in ways that may have looked familiar to a Victorian gentleman when interacting with servants. However, robots potentially also subvert these protective strategies, when sensors and their capacity are hidden, or worse, design choices create the potential to actively mislead people who interact with robots about their potential. We suggested a transparency duty to mitigate this risk. However, robots are also "vulnerable" and can become a danger to others, when their sensors are interfered with, disabled or fed misleading information. This is in particular the case when the robot uses (also) machine learning. Robots, to operate safely, require a cooperative environment where certain data and information is made available to them, something that transparency of their sensory capacities can jeopardise. We suggested therefore a quid pro quo: robots *should* be built with their sensory capacity openly displayed, in return, people interacting with machines *should* volunteer information the root needs for safe operation, even though this may intrude into their (privacy or intellectual property) rights position. Furthermore, they *should not* manipulate the sensors of the robot if this can foreseeably cause a danger – "gaslighting" a robot.

3. Closing the circle: fair data handling and benefit sharing

In our suggested quid pro quo, a final element is missing, and this brings us back to the duties of the robot designer or operator. The discussion so far discussed problems caused by humans who withhold/distort/manipulate data that a robot needs to operate safely. We argued that the ethically mandated transparency of sensory abilities may need to be balanced against the likelihood that a degree of non-cooperation will take place. However, we also argued for an ethical principle of information disclosure that encourages people to contribute to safe robots by making available information that under data protection or IP rules, they could also chose to withhold – in the most obvious example, refrain from interference with a robot's sensors. This however means that we also face ethical design choices when people cooperate and volunteer information that they are not legally obligated to provide, but chose to provide/do not inhibit out of a sense of civic duty. In this case, the robot owner/designer benefits from the voluntary disclosure of information (as opposed to the "sneakily obtained information due to deceptive sensor design"). We should then ask if this benefit triggers in turn ethical or legal obligations, and if it affect the status of any *output* the robot produces. In medical research for instance, people volunteer information about

themselves that they could easily withhold. Good governance systems respond to this by putting additional demands on the "fairness" of the use of this data, which can take e.g. the form of benefit sharing (see e.g. Chadwick and Berg 2001). For volunteering information that people were not obligated to provide, at the very least they should be protected from any negative consequences of this disclosure, and ideally should benefit from it. In concrete terms, information disclosed in a medical trial must not be used to raise the insurance premium of the test subjects, and ideally they should get symbolic recognition (if so wanted) or material benefits such as a discount when a drug is developed through their help. We introduced briefly an example from robotics above that can help us to analogize the situation. An industrial robot is introduced at the workplace, and through observing human workers learns how to safely interact with them. We argued that this robot should not mislead the workers as to its sensory capacities – not e.g. have an eye-like structure when in reality it uses acoustic and heat sensors placed at a different part of the machine. This allows the workers to chose privacy preserving strategies when they feel the need, such as turning their back to the machine when having an unauthorized cigarette break that a heat sensor would detect. In return, they *ought* to be willing to let themselves be observed when it matters, in particular, when the robot learns form their movement on how to best perform their task, as a requirement for safe interaction. This even though it took them years to optimize their movements, and it is this skill that makes them valuable as workers. In return though, and following the "benefit sharing" analogy, information obtained for the safe operation of the robot must not be used for the purpose of replacing the workers eventually by that machine (in data protection terms, "purpose limitation" of data even if, as here, it is doubtful if the information in question is personal information for DP purposes), and ideally their contribution should be acknowledged and rewarded if the robot becomes a commercial success (even though the contribution they made in our scenario is not one protected by IP, the are treated as if the robot's newly acquired knowledge is a "derivative work" of their know how)

Much of this of course goes beyond the remit of the robot developer, and is more a question of the ethical deployment of robots in commercial or social settings. However, it does affect certain design choices too. For instance, it could mean to securely delete data other than personal identifiable data as soon as the original purpose, safe navigation of a space, was achieved. It could also mean that not only robots need to be identifiable as robots and their sensors as sensors, but robot generated output also must be identifiable as machine, not human generated. An appropriate legal regime can assist in this complex quid pro quo, and balancing of competing interests. Copyright law might for instance impose relevant constraints on the way in which machines communicate the status of their outputs: in jurisdictions that do not protect computer generated works, a "this text was generated by algorithm that learned from observing John and Jane Doe, feel free to share" might be required. In such a setting, John Doe's input is "treated fairly" by acknowledging his contribution short of authorship, and by making the outcome freely available to everybody including them.

The overarching theme of this intervention, therefore, is ultimately one of algorithmic transparency: legal and ethical duties influence when, and how, robots should disclose their sensory capacities. Ideally, robots will be transparent about their sensory capacity. Once they do this, their environment has choices – to cooperate or not to

cooperate. Where non-cooperation causes harm, there might be wider social discussions to be had under what conditions it is more beneficial to allow clandestine gathering of (some) data to have safer machines, or if we need to create new legal duties to actively assist, or at east not to misdirect, a robot. Where cooperation beyond the legally required produces value, a discussion needs to be had how to account for this in an equitable way, e.g. imposing potently another transparency duty such as a "made by robot with assistance of" label, or even a form of benefit sharing that we encounter routinely in medical research. Pro-ethical design will be able to accommodate and cater for these choices, by limiting for instance data gathering and storage beyond what Data protection law requires anyway for personal identifiable data. While these are mainly legal issues, for the question of ethical (and law compliant) design, the developers need also to be able to anticipate what type of interaction to expect given the normative overall framework, including duties owed towards the robot owner or manufacturer by third parties The developer's ethical obligations under the Principles do not exist in a social vacuum, but have to be informed also by a dynamic understanding and anticipation of the wider normative environment.

Arendt, H. (1959). *The human condition: a study of the central dilemmas facing modern man*. Doubleday, New York

Chadwick, R., & Berg, K. (2001). Solidarity and equity: new ethical frameworks for genetic databases. *Nature Reviews Genetics*, *2*(4), 318-321.

DeCew, J. W. (1997). *In pursuit of privacy: Law, ethics, and the rise of technology*. Cornell University Press

Fradella, H. F., Morrow, W. J., Fischer, R. G., & Ireland, C. (2010). Quantifying Katz: Empirically Measuring Reasonable Expectations of Privacy in the Fourth Amendment Context. *Am. J. Crim. L.*, *38*, 289 – 374

Hamill, L. (2006). Controlling smart devices in the home. *The Information Society*, *22*(4), 241-249.

Kang, J. (1998). *Information Privacy in Cyberspace Transactions*, 50 Stan. L. Rev. p. 1193 -129

MacDorman, K. F., Minato, T., Shimada, M., Itakura, S., Cowley, S., & Ishiguro, H. (2005, July). Assessing human likeness by eye contact in an android testbed. In *Proceedings of the XXVII annual meeting of the cognitive science society.* Mahwah: Lawrence Erlbaum Associates, 21-23

Sag, M. (2009). Copyright and Copy-Reliant Technology. *Nw. UL Rev.*, *103*, 1607-1682

Xu, T. L., Zhang, H., & Yu, C. (2016). See You See Me: The Role of Eye Contact in Multimodal Human-Robot Interaction. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, *6*(1), 2

Yonezawa, Tomoko, et al. (2007). Gaze-communicative behavior of stuffed-toy robot with joint attention and eye contact based on ambient gaze-tracking. In: Proceedings of the 9th International Conference on Multimodal Interfaces. ACM, New York 140 - 145