2017 IEEE International Conference on Cloud Engineering

Eyes of Things

Noelia Vallez*, Jose L. Espinosa-Aranda*, Jose M. Rico-Saavedra*, Javier Parra-Patino*, Oscar Deniz*,
Alain Pagani†, Stephan Krauss†, Ruben Reiser†, Didier Stricker†, David Moloney‡, Alireza Dehghani‡,
Aubrey Dunne‡, Dexmont Pena‡, Martin Waeny§, Pedro Santos§, Matteo Sorci¶, Tim Llewellynn¶,
Christian Fedorczak‖, Thierry Larmoire‖, Marco Herbst**, Andre Seirafi†† and Kasra Seirafi††
*VISILAB Group, University of Castilla-La Mancha, Avda Camilo Jose Cela s/n, Ciudad Real, Spain
†DFKI, Augmented Vision Research Group, Tripstaddterstr. 122, 67663 Kaiserslautern, Germany
‡Movidius, 1st Floor, O'Connell Bridge House, D'Olier Street, Dublin 2, Ireland
§Awaiba, Madeira Tecnopolo, 9020-105 Funchal, Madeira - Portugal
¶nViso SA, PSE-D, Site EPFL, CH-1015 Lausanne, Switzerland
‖THALES Communications & Security, 4 Avenue des Louvresses, 92230 Gennevilliers, France
**Evercam, 6-7 Granby Row, Dublin 1, D01 FW20, Ireland.
††Fluxguide, Burggasse 7-9/9, 1070 Vienna, Austria

***Abstract*—Responsible Research and Innovation (RRI) is an approach that anticipates and assesses potential implications and societal expectations with regard to research and innovation, with the aim to foster the design of inclusive and sustainable research and innovation. While RRI includes many aspects, in certain types of projects ethics and particularly privacy, is arguably the most sensitive topic. The objective in Horizon 2020 innovation project Eyes of Things (EoT) is to build a small high-performance, low-power, computer vision platform (similar to a smart camera) that can work independently and also embedded into all types of artefacts. In this paper, we describe the actions taken within the project related to ethics and privacy. A privacy-by-design approach has been followed, and work continues now in four platform demonstrators.**

I. Introduction

Eyes of Things [1] (Grant n. 643924) is an Innovation Project funded by the European Commission within the Horizon 2020 Framework Programme for Research and Innovation. It started January 1st, 2015 and will end December 31st, 2017. The objective in EoT is to build an optimized core vision platform that can work independently and also embedded into all types of artefacts. In practice, this means a small (can be wearable), low-power (can be battery operated) flexible 'intelligent camera'.

EoT is motivated by the phenomenal advances and confluence of camera technologies, efficient low-power processing and computer vision. Camera technologies and low-power processing have gone parallel to the success of mobile devices. Computer vision, i.e. the automatic analysis of camera images to extract meaning or otherwise useful information, was once restricted to quality inspection in factory floors. However, the discipline is rapidly progressing to the point of leading to applications almost everywhere. A very first example of the real potential of computer vision is in the Microsoft XBOX Kinect vision sensor, which is still the fastest-selling consumer electronics device ever. In smartphone cameras, vision techniques such as face detection, smile detection and panorama generation are now commonplace. There are many more examples: image search, license plate recognition, video stabilization in YouTube, Facebook's facial recognition for photo tagging, etc. The advent and progress in deep learning, UAVs (drones) and augmented/virtual reality also contribute to this rosy scenario for computer vision. In this context, EoT aims at being a flexible platform for OEMs to develop computer vision-based products and services in a shorter time.

EoT is being developed by a Consortium of European partners: one University (coordinator), one research centre, 4 SMEs and two multinationals. Roughly the first half of the project (up to month 21 of 36) is devoted to the development of the hardware and software platform. During the second half of the project the platform is to be put to use in four demo scenarios. At the time of writing, the platform is almost finished. Nearly 400K lines of software code are available in the form of libraries and sample applications. This includes libraries for general-purpose image processing and computer vision, QR code recognition, Python scripting language (besides C language), deep learning inference, video streaming, robot control, audio input and output, efficient wireless communication, etc. As for the hardware, the main components are: a) highly efficient Myriad 2 processor, b) tiny low-power camera and c) low-power WiFi connectivity. Flash memory and a SD card are available for storage. Figure 1 shows the aspect of the final factor-form board, which is 7x5cm.

EoT is an example Cyber-Physical System (CPS). CPSs can be seen as an evolution of embedded systems in an Internet of Things paradigm: the interconnectivity features are enhanced, devices have extremely reduced size and cost and in general they can be found in higher number almost everywhere. While CPSs promise many benefits for society, the effects of newly introduced technologies can never be completely predicted. In particular, CPS may collect vast amounts of data, and this poses several privacy questions [2].

This paper summarizes the work done so far in EoT in terms of RRI. We want to emphasize that 1) this is ongoing work and 2) the effort made in RRI is limited to the resources allocated within the project for this aspect. Advice and guidance was

978-1-5090-5817-4/17 $31.00 © 2017 IEEE
DOI 10.1109/IC2E.2017.44

IEEE computer society

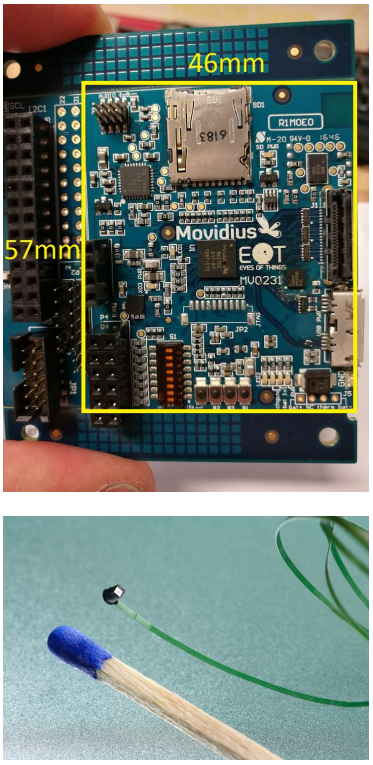

Fig. 1. Top: EoT board. Bottom: camera.

given in by the EoT Ethics Board as well as by EU FP7 project *Surveillance: Ethical Issues, Legal Limitations, and Efficiency* (SURVEILLE) [3]. The rest of the paper is organised as follows. Section II describes the privacy-by-design features implemented in the platform. Section III discusses the ongoing work on the four demonstrators being developed using the platform. Finally, the main conclusions are drawn in Section IV

II. Privacy-by-design

Privacy by Design [4] is an approach to systems engineering which takes privacy into account throughout the engineering process. Essentially, the concept refers to taking human values into account even when this may counter technical optimizations or consume development resources. Privacy by design is an approach to privacy by embedding it into the design process of the new technologies.

EoT is a device which is intended to capture and process images. Images will, more often than not, depict people in a way that would allow them to be recognized, thus representing personal data. It is therefore crucial that we implement mechanisms that allow preserving privacy and security. This was understood by the Consortium even before the project started.

In this section we describe the major technical efforts made to provide privacy and security to the EoT platform. As soon as hardware specifications were clear, it was observed that the two main points of access to (personal) data by external agents were the WiFi communication and the integrated SD card. Therefore, most of our efforts have been focused on providing means to secure the use of those elements. Some of these security features are embedded into the design and cannot be avoided, while others are optional and are on by default.

A. Default boot access-point

By default the EoT device does not connect to any WiFi on boot. Rather, it creates a WiFi to which only one other device can connect initially, assuming it knows the password. This is a security mechanism intended to restrict connection to a home WiFi and Internet. With this mechanism, the default use would be that of a client connecting from a laptop or tablet. Connection can only be established if the client knows the access-point password, which is stored in the device flash memory and is unique to each manufactured device.

While this is the default behaviour, an authenticated client can configure the device to make it connect to an existing WiFi.

B. WPA2

Wi-Fi Protected Access (WPA) is a security protocol developed by the Wi-Fi Alliance to secure wireless computer networks. The more advanced WPA2 replaced WPA in 2004. WPA2 is the security protocol implemented in EoT for all WiFi communications.

WPA2 includes encryption and client authentication. In fact, one of the most significant changes between WPA and WPA2 was the mandatory use of AES algorithms and the introduction of CCMP (Counter Cipher Mode with Block Chaining Message Authentication Code Protocol) as a replacement for TKIP (still preserved in WPA2 as a fallback system and for interoperability with WPA. The use of a long random password makes WPA2 virtually uncrackable.

C. Encrypted SD Card

The EoT includes a Secure Digital (SD) Card to store data. Secure Digital cards are used in many consumer electronic devices, and have become a widespread means of storing several gigabytes of data in a small size. The card is physically accessible. To remove the card, all it takes is to push the card slightly to eject it. Moreover, these SD cards can be read in modern laptops, which come with an SD card slot. This makes the presence of an SD card in EoT specially sentitive. While some SD cards have a small sliding tab to make it write-protected, this mechanism is obviously not sufficient.

In order to ensure security of the card, an encrypted filesystem has been implemented. This means that all card's contents are stored in an encrypted form. File-level encryption was implemented using AES-CTR with 128 bit key length. The Advanced Encryption Standard (AES) is a specification for the encryption of electronic data established by the U.S. National Institute of Standards and Technology (NIST) in 2001. AES has been adopted by the U.S. government and is now used worldwide. It is the first (and only) publicly accessible cipher

approved by the National Security Agency (NSA) for top secret information. AES is a symmetric-key algorithm, meaning the same key is used for both encrypting and decrypting the data. Needless to say, this encrypted filesystem introduces additional overhead to all card operations.

D. File shredding

Data erasure (also called data clearing or data wiping) is a software-based method of overwriting the data that aims to completely destroy all electronic data residing on a hard disk drive or other digital media. Permanent data erasure goes beyond basic file deletion commands, which only remove direct pointers to the data disk sectors and make the data recovery possible with specific software tools.

Old media types are especially sensitive. As an example, data on floppy disks can sometimes be recovered by forensic analysis even after the disks have been overwritten. However, it is widely acknowledged that with modern, non-magnetic media one overwrite is all that is generally required.

Apart from the basic 'delete file' function in the SD card filesystem, in EoT we have implemented a 'shred file' functionality that overwrites the file with 0s prior to its deletion. Needless to say, this operation takes much longer than the basic 'delete file'.

At the suggestion of the Ethics Board, both the encrypted filesystem and the shred operation were made the default operating modes.

E. Google Cloud Vision API

The Google Cloud Vision API is a cloud-based service for image analysis. To use the service, a device uploads an image to the corresponding Internet servers and metadata is returned with information about the image. Currently, the service supports:

- Detecting broad sets of objects in the images, from flowers, animals, or transportation to thousands of other object categories commonly found within images.
- Detecting inappropriate content
- Facial detection and analysis emotional facial attributes of people in the images, like joy, sorrow, and anger. Facial recognition is not supported.
- Recognising text within the images, along with automatic language identification.
- Logo Detection: Detect popular product logos within an image.
- Landmark Detection: Detect popular natural and man-made structures within an image.

While it was not in the original project plan (the API it was only released in beta form in late 2015), the Google Cloud Vision API greatly extends EoTs capabilities. Note that other similarly powerful cloud-based APIs provided by other vendors exist. Access to the Google Cloud Vision API in EoT has been facilitated by a software module developed within the project. However, we have to emphasize that this is an entirely optional module. As in most other code developed in EoT, its use depends on the final application.

Access to the Google Cloud Vision API is always done using Secure Sockets Layer (SSL). SSL is designed to provide two principles: privacy and authentication. SSL is the underlying secure communication layer of the *https* protocol widely used on the Internet. In SSL, privacy is achieved by encryption/decryption and authentication is achieved by signature/verification. Besides, as per Terms of Service, Google ensures protection and fair use of images.

The Google Cloud Vision API is a commercial service (for more than 1000 queries/month). In order to use it, the programmer has to create a billing account with Google. In EoT we have not used the commercial tiers of the service. Besides, as part of a collaboration with Google Research, the consortium obtained additional levels of use to test the service internally. Any prospective programmer (outside of the EoT Consortium) will need to create his/her own billing account. The holder of the billing account is responsible for accepting Google's terms of service.

The EoT code to facilitate the use of the Google Cloud Vision API in EoT is essentially made up of a core library (which provides access to the API) and a number of examples. The examples are applications that use the library. The examples capture images from the camera and use functions in the library to securely send them to the service provided by Google. After sending the images to the service the images are deleted from the EoT device. Images are never stored in any non-volatile memory in the EoT device. The text strings returned from the Vision API are received in the EoT device and shown to the user in different ways. One example application is an OCR (text recognition) demo. In this example, the application captures images. The user is expected to place a sheet of paper with a 3-digit number in front of the camera of the EoT device. The image is sent to the Vision API, which returns the 3-digit number as a string. The latter is then used to feed a TTS synthesizer software running in EoT that uses the EoT audio output to play the pre-recorded sounds of the 3 digits. This and similar examples are the only way in which this service is used within EoT, and always by researchers of the EoT Consortium. These examples will not be used to capture people other than those EoT researchers doing the demonstrations. For these demonstration purposes the EoT Consortium is acting as data controllers, albeit always ensuring the principles of transparency, legitimate purpose and proportionality defined in Directive 95/46/EC on the protection of individuals with regard to the processing of personal data and on the free movement of such data. In this use of the Vision API, Google acts as a processor of data, and as such Google provides a level of security and privacy appropriate to their access to data and the scope of the services they are engaged to provide. Notwithstanding, the transfer of personal data outside of the EU is affected, according to Google's terms of service, by the Safe Harbor Principles, which have been actually ruled as invalid by the European Court of Justice in 2015. The legal situation is currently uncertain in this respect. This means that, until the new General Data Protection Regulation enters into force (25 May 2018) or a new

transatlantic agreement on privacy is reached, this processing must not be considered valid in general. Therefore, the use of this API in EoT must not go beyond demonstration purposes by EoT researchers.

If, at any point within the project lifetime, third parties are given an EoT unit and this software (as part of an Early Adopter Program, for example), they will be required to a) sign an agreement in which they bind themselves to the same type of demonstration use as described above and b) to use their own billing account. Further, they will explicitly acknowledge the extant data transfer scenario described above, as long as they act as controllers processing personal data within the EU.

Note that this scenario already imposes restrictions on our work within the project. The powerful emotion recognition functionality of the Vision API, for example, could have been used in the doll demonstrator (see below). The situation, however, precludes the use of the Vision API in this demonstrator.

F. Other features

An image of an identifiable individual is considered personal data. Therefore, face blurring or pixelation will be an important capability that must be present in EoT. In fact, given the potential of the EoT device, which in some respects such as size and portability go beyond the state of the art in cameras and makes surreptitious use a real possibility, the capability must be implemented and it should be, whenever possible, the default mode. Note that it is not necessary to blur the whole image (which would make the image useless) but only identifiable individuals.

Face detection was implemented in EoT. However, face detection algorithms work with upright faces. If the head is rotated (beyond 15-20 degrees approximately) it will not be detected. Analogously, if the head is upright but the EoT device is rotated (which is quite likely given that it will be very small and mobile) then the face will not be detected either. Therefore we implemented a mechanism to detect faces even in those cases. The algorithm is based on the IMU included in the EoT device. The input image is first rotated according to the angle read from the IMU, and then a standard upright face detector is applied. This can be seen in the Figure 2, which shows face detection and blurring when the device was being rotated. Note how black borders appear due to the rotation applied to the image.

III. EoT DEMONSTRATORS

The EoT platform is to be demonstrated with four example applications. These four demonstrators are only meant as prototypes to show specific capabilities implemented in the platform. In some of them there are possibilities for commercial exploitation, although that is out of the scope of the project.

Note that all demonstrators are actually made up of two software elements: the application running in the EoT device and an App running on the user's smartphone, tablet or computer. The latter is used to install the EoT application and configure it.

Fig. 2. Face blurring. Note that the figure is showing the rotation-invariant face detection capability of the platform, whereby the device's IMU is used to counter tilt so that faces can still be detected with the standard upright face detector. This feature was added to consider the possible uses as a wearable or inside a toy.

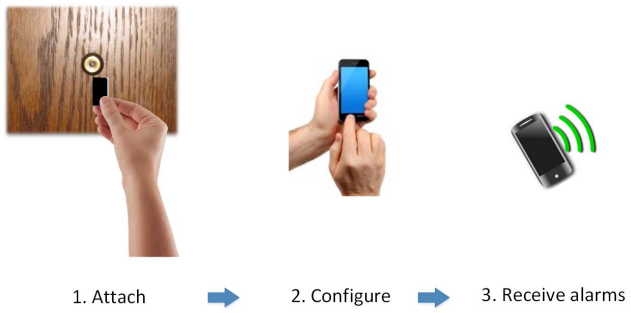

Fig. 3. Peephole surveillance.

A. Demonstrator 1: Peephole surveillance

The context of this demonstrator is the following: video surveillance is a growing market, from big systems (city scale) to smaller ones (home security). This latest one is growing exponentially and requires low-cost and easy-to-install, easy-to-use systems. In this context, the first demonstrator will develop a peephole surveillance system. Before leaving home, the user will attach the EoT-based device to the peephole. The device will continuously monitor for presence and/or suspicious activity at the door, sending alarms and pictures via Internet (assuming home Wi-Fi is available), see Figure 3. The device will not need cables, since it will function with its own rechargeable battery. Tampering detection will be also implemented (i.e. an attempt to cover the peephole) and will also generate an alarm.

Even though we intend to focus on 'safe' experimental laboratory scenarios within the Consortium, it is important that we consider and follow the recommendations made by the project's Ethics Board. In particular, the board made the interesting observation that the relatively low resolution of the camera (which was in our case imposed by the need to reduce power, and viewed from our side as a technical limitation) may be in fact an advantage. This demonstrator is only intended to capture images of people right in front of the door, and not

of those farther away (possibly neighbours). Low resolution can provide relatively good images of the person in front of the door, but would make people farther away unidentifiable. So, instead of relying on informed consent (which would be difficult to implement in practice, though still possible), technical work can in fact help achieve anonymity, at least in three possible ways:

- Lowering image capture resolution to the point of making people not standing right in front of the door unidentifiable. Resolution could be a configuration parameter in the application
- Blurring or masking all faces that are detected by the application and have a size smaller than a given, again configurable, threshold. This would anonymize all faces except those of people standing at the door
- Implementing a foveal transformation of the image, whereby only the central zone remains sharp

Options 2 and 3 could be the most effective in practice. A combination of methods is also possible.

On the other hand, the communication of data, particularly beyond the European Union, is currently problematic. The image of the person in front of the door will have to be securely transmitted. Communication beyond the European Union is more difficult to control from a technical point of view and may be dependent on decisions that are only taken at a potential commercialization phase (like the selection/contracting of the server that receives images from the home setup and forward the images to the user's smartphone). Still, note that even if the server is secure and is in the EU, the user may eventually be outside of the EU, so the terms and conditions of use of this application would need to state that use is only permitted within the EU.

B. Demonstrator 2: Museum audio tour

This demonstrator envisions ways in which the EoT solution can be used for interpretation and learning at museums, cultural heritage sites, exhibitions, and similar venues. The prototypic outcome of the demonstrator will be a visitor service system which automatically detects exhibits and then provides information about them to the user. Current museum interpretation is either based on (a) classical audio information (audio guide systems) or (b) on more elaborated, rich content/interactive systems (touch device apps). We will show that the EoT core system is able to cover both paradigms and to open up a totally new solution for museum experts and museum visitors. The main use case (The invisible museum guide, see Figure 4) is based on automatic exhibit detection which results in an unprecedented, context sensitive information and interaction system for museum visitors.

Within the project's lifetime, pilot tests will be conducted in real museums, though only during the closing hours or within a dedicated area in the museum closed to public. As far as we understand, these pilot tests of the demonstrator do not raise ethics issues. As for the legal implications in a real, commercial implementation, no issues are foreseen as long as the device does not record images of the artwork.

Fig. 4. Museum audio guide demonstrator.

C. Demonstrator 3: Wearable camera

This demonstrator describes a body worn camera powered by EoT. The device would be similar in design to existing *life-recording* body-worn cameras, and would benefit from the advanced capabilities of the EoT platform (instead of recording everything, record meaningful or interesting images only). There will also be a cloud element to the demonstrator that provides the user interface to the device, the captured media and enhanced apps built thereon.

All of our experiments/trials in this demonstrator will be performed by the researchers in the Consortium, and always in indoor conditions (i.e. no use outdoors). This is a safe scenario, with the minimal necessary commitment to demonstrate the technology. In any case, in order to ensure privacy, there will be at least two modes intrinsically built into the demonstrator. In the default mode, detected faces will be blurred to anonymise them. This will happen in the EoT device and right after image capture. The user interface of this demo application will present the user with the two possibilities (privacy on, privacy off), with 'privacy on' being the default. Although this demonstrator may not by itself turn into a commercial product, it will be in any case disseminated as a potential application. Our aim in doing this is also to raise awareness in relation to the question of privacy. In fact, EoT may well be *The world's only ethical, privacy-compliant camera*.

D. Demonstrator 4: Doll with emotion recognition

This demonstrator will embed an EoT device inside a doll's head (or torso). Facial emotion recognition will be implemented so that the doll can assess the girl's emotional display and react accordingly (with audio feedback). Two scenarios had been considered in this demonstrator targeting a usage of the interactive doll in playful situation in the first scenario and as a therapeutic tool in the second.

The use of a camera or other sensors in a toy used by children raises many ethical and legal issues. A well-known

example was the Hello Barbie produced by toymaker Mattel. The doll has a mode in which everything that a child says is transmitted to cloud servers. These capabilities have raised concern and protests by advocate groups. Moreover, the doll has also become an example of security vulnerabilities (which have, unfortunately, proven unavoidable in IoT devices so far). Hackers have already managed to access the doll to get data out of it, including account IDs and MP3 files. More recently, in a development taken up by the BEUC (representing consumers at EU level), the Norwegian Consumer Council has found that toys recording sounds (presumably, similar arguments can be made for video recordings too) violate EU law [5]. The products are being challenged not just directly on data protection grounds but also consumer protection and safety. Even in our case, in which all processing takes place in the device, the possibility remains that someone hacks the device so as to replace the legitimate software running on it with other malicious software. The latter could clearly invade children's privacy and allow for criminal uses. This is potentially a toy that can record and steal images of the child, even when the child is not using the toy. Such software can be also used as 'ransomware'.

The consortium acknowledges that giving up on such scenario is not in accordance with the 'positive-sum' principle of privacy-by-design [6]. However, as far as we know there are no security features that we can add to avoid the real possibility that security is breached either digitally or by physical means, with catastrophic consequences. The level of security must be always proportionate to the potential harm. In this case, there is no sufficient security in the basic elements of the technology used (and nothing that we could have added to them to remedy that fact) to counterbalance the potential harm.

This is an interesting case in which the application of the positive-sum principle can require resources or technology beyond those available. In such cases, a higher principle of responsibility must prevail.

Therefore, in an exercise of responsibility the consortium will focus on the second scenario, which already involves a number of challenges that will have to be addressed. The Ethics Board concurred that this is a 'safe scenario' to show the capabilities. Note, for example, that physical access to the devices would be clearly restricted in this setting. The partner involved in this demonstrator (nVISO) has experience in this regard and actively collaborates with a hospital-based team of researchers and experts.

IV. Conclusion

This paper summarizes the approach taken within EoT to help safeguard fundamental rights. This includes technical decisions both at the platform development level and within the demonstrators. Work in EoT is still ongoing, mainly to fully develop the four demonstrators. All throughout, close attention will be also paid to Article 29 Working Party discussions. This group, which is made up of the independent regulatory authorities from the 28 EU Member States, discusses issues related to data protection and clarify the meaning of the legislation or its application to new technologies. Their opinions are not formally binding although they are nonetheless persuasive and most valuable. So far we have identified some opinions that might be relevant to the eventual commercialisation of the demonstrators:

- On smart devices and 'Apps' [7]. This opinion focus on the general lack of transparency and awareness of App users, invalid consent mechanisms and a trend towards 'data maximisation and elasticity of data processing purposes'. Recommendations in this opinion are aimed at App developers, since they have the greatest control over the precise manner in which the processing is undertaken or information presented within the App.
- With regard to cloud platforms and the C-SIG Code of Conduct [8].
- Legitimate interests of the data controller (this can justify processing rather than having to rely on consent) [9].

Eyes of Things is a project that exemplifies the current global trend towards cognitive applications. It is also an example of the IoT paradigm whereby a myriad of sensors gather data from the world and us. Besides the technical challenges and the potential applications, from the point of view of ethics and privacy it already touches on sensitive issues that will have to be broadly addressed in the near future.

Acknowledgment

The authors would like to thank the EoT Ethics Board: Petra Ahrweiler, Daniel Neyland and Lorna Woods for the fruitful discussions and recommendations. The authors would like to thank the advice provided by EU FP7 project SURVEILLE. This work has received funding from the European Union's Horizon 2020 Research and Innovation programme under grant agreement No 643924.

References

[1] EoT Project. Last accessed: 14/12/2016. [Online]. Available: http://eyesofthings.eu

[2] European Parliamentary Research Service, "Ethical Aspects of Cyber-Physical Systems. Scientific Foresight Study," Science and Technology Options Assessment Panel, Tech. Rep. PE 563.501, June 2016.

[3] SURVEILLE. Last accessed: 14/12/2016. [Online]. Available: https://surveille.eui.eu/

[4] P. Hustinx, "Privacy by design: delivering the promises," *Identity in the Information Society*, vol. 3, no. 2, pp. 253–255, 2010. [Online]. Available: http://dx.doi.org/10.1007/s12394-010-0061-z

[5] Connected toys violate European consumer law. Last accessed: 14/12/2016. [Online]. Available: http://www.forbrukerradet.no/siste-nytt/connected-toys-violate-consumer-laws

[6] A. Cavoukian, "Privacy by Design: The 7 Foundational Principles," Mai 2010, revised: October 2010.

[7] Opinion 02/2013 on apps on smart devices. Last accessed: 14/12/2016. [Online]. Available: http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2013/wp202_en.pdf

[8] Opinion 02/2015 on C-SIG Code of Conduct on Cloud Computing. Last accessed: 14/12/2016. [Online]. Available: http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2015/wp232_en.pdf

[9] Opinion 06/2014 on the notion of legitimate interests of the data controller under Article 7 of Directive 95/46/EC. Last accessed: 14/12/2016. [Online]. Available: http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2014/wp217_en.pdf