Ethics in Robotics Research

CERNA Mission and Context

© PHOTOCREDIT

By Alexei Grinbaum, Raja Chatila, Laurence Devillers,
Jean-Gabriel Ganascia, Catherine Tessier, and Max Dauchet

This article summarizes the recommendations concerning robotics as issued by the Commission for the Ethics of Research in Information Sciences and Technologies (CERNA), the French advisory commission for the ethics of information and communication technology (ICT) research. Robotics has numerous applications in which its role can be overwhelming and may lead to unexpected consequences. In this rapidly evolving technological environment, CERNA does not set novel ethical standards but seeks to make ethical deliberation inseparable from scientific activity. Additionally, it provides tools and guidance for researchers and research institutions.

Digital Object Identifier 10.1109/MRA.2016.2611586
Date of publication: 18 January 2017

Handling Ethics

A broad reflection on the ethical consequences of automation, robotics, and artificial intelligence began simultaneously with their emergence as research disciplines. Norbert Wiener, one of the founders of cybernetics, raised a number of questions about the transformation of human society by these nascent technologies [1]. In recent work [2], [3], emphasis often has been placed on the possible catastrophic consequences of information technology and artificial intelligence. Ethical reflection in the context of modern robotics dates back to the turn of the century, when progress in this field reached a level of maturity that motivated several researchers to raise concrete ethical concerns that were grounded in real applications. In 2002, a research atelier was funded by the European Robotics Research Network that drafted the Roboethics Roadmap [4]. In 2004, the IEEE Robotics and Automation Society established a technical group on roboethics. A special issue of *IEEE Robotics and Automation Magazine* was devoted

to roboethics in March 2011 [5]. The European Union funded several projects on this issue, including ETHICBOTS [6] and Robolaw [7]. In 2016, the IEEE launched a major initiative on the ethics of autonomous systems [8].

Among all applications of robotics, the military use of robots is the most controversial domain, which, for the research community and the public, has contributed to an ever increasing awareness that an ethical study of robot conduct has now become mandatory. Ronald Arkin authored one of the pioneering works in the area [9], in which he explicitly addressed ethical decision making for autonomous robots within the context of the battlefield. Ethical, legal, and societal issues that have been raised by robotics and artificial intelligence can now be found among the concerns of governments and corporations worldwide. The use of autonomous weapons is under discussion at the United Nations. The European Parliament has created a working group on robotics and artificial intelligence [10].

CERNA was instated in 2012, and, in November 2014, it published a report that specifically addressed the ethics of robotics research [11]. This article is a summary of that report's key chapters. CERNA recommendations have a limited scope: they only concern researchers and engineers at the stages of preparing, designing, and implementing their research project.

Robotics and, more broadly, information science and technology affect the lives of all citizens via their influence on the economy and employment. This could lead to massive, albeit unexpected, consequences. Currently, for example, we are witnessing a boom in the production and ownership of civilian drones, which may, by analogy with the automobile, cause a change in the way that human spaces are used and how they look. Among the various effects produced by cars, one that was difficult to imagine at an early stage was the existence (and the very concept) of traffic jams; we do not know yet whether drones or autonomous vehicles will prove to be equally instrumental in producing a new social reality. This essential uncertainty is the reason why we believe it is inappropriate to set novel ethical standards in such a rapidly evolving technological environment. If new rules were decreed based on current moral judgment, they might quickly become obsolete. Instead, we seek to empower scientists to embrace ethical issues in a way that is inseparable from their research.

CERNA structured its reflection on the ethics of robotics research into three areas: autonomy and decision making; the imitation of living beings, including affective and social interactions with humans; and robot-assisted therapy and human-robotic augmentation and enhancement. This integration of ethical thinking and research should have both a collective and an individual dimension. To this end, CERNA general recommendations include the creation of operational ethics committees at ICT research institutions and dedicated support activities for raising researchers' awareness. While running a project, scientists should consult such operational ethics committees on emerging controversial issues that

require collective deliberation. CERNA also recommends that research institutions set up interdisciplinary research projects to address the larger body of legal and socioeconomic issues in robotics.

Robots in Society

Integration into Society

Researchers should protect the systems they design from undesired effects. This is particularly important in robotics because robots are increasingly endowed with autonomy. Among machines that count as service robots, most possess limited autonomy and are designed for specialized tasks in the context of limited use. This is true of all existing professional service robots. Today, the bulk of robots sold to private individuals largely comprises automated cleaning devices and toys. Extensive research efforts exist on robots in open environments that require greater autonomy, including an advanced capacity for interacting with humans as well as increased learning capabilities. Researchers should take into account the required level of trust in a robot, its capabilities and limits, and the capabilities and limits of the partnership it forms with the user. They should analyze how a robot can be controlled and traced, i.e., how its behavior can be reported or understood.

Humanoid robots are constantly improving. They can be programmed to dance, take hold of objects, imitate gestures, or play football. They communicate and interact in a rudimentary manner through gestures or speech, and some are able to mimic emotions. By imitating a living being and through emotional interaction, a robot may be instrumental in blurring the boundary between machines and humans. It may also play with emotions in a completely novel way. Such robots are currently used for research, including medical studies, or for educational purposes. Other robots are used as personal assistants, in particular, for the elderly. Android robots may give rise to excessive fears and hopes due to their shape. This is often amplified by hype, ideologies, or beliefs. In this domain in particular, responsible roboticists should remain aware of the degree to which the precise state of science and technology differs from its image among the general public. They should seek to enlighten the debate through measured and scientifically informed communication.

A few research labs seek to achieve the development of a robot with a high degree of emulated human likeness or behavior. Beyond technological prowess, the problem of necessity and usefulness of robot-human likeness should be raised, and a cross-disciplinary assessment of its effects conducted, in particular, for robots designed to interact with children or vulnerable adults.

Current developments support the idea that robots may soon be operating in society alongside humans. Such robots should no longer be considered as standalone technical objects; rather, they become sociotechnical systems with varying degrees of autonomy and integration. Although a robot may remain clearly distinguishable

from a living being, the imitation of a particular natural trait (e.g., human likeness or a human-like voice) gives it a special place in social interaction. This problem of status is important as well as culture dependent: in antiquity, a statue with colored and moving eyes or an artifact capable of speech was taken to be alive and often provoked reactions similar to those induced by living beings. Ethical and legal issues do not only concern a technical object per se, its design or reliability; they also include interaction among robots, humans, and society.

It is important to spell out the goals of a project in robotics, particularly in social robotics, at the design stage: For whom is the robot intended? In what areas can it be used? Is the project likely to have a major impact on the lives or well-being of robot users? Which stakeholders are involved? Such a preliminary analysis is essential as it supplies the deliberative process with explicit choices to be made. Researchers should carefully document the designed system, seek to explain its capabilities and limits, and remain aware of the effects of hype in public communication.

Respect for Privacy

Some robots, such as caretaker and surveillance robots, personal assistant robots, or drones, are capable of collecting personal data (photos, videos, voice recordings, physiological parameters, geolocation, etc.). Their deployment raises issues related to the privacy and protection of personal data. While it is not possible at the design stage to protect a robot from the inappropriate or illegal use of the data that it collects, researchers must nevertheless remain watchful to ensure that the robotic system facilitates monitoring and control of data in accordance with the existing regulation.

Legal Aspects

The existing legal frameworks cover a large variety of legal issues in robotics. However, there is an ongoing debate about new standards, which extends to the entire domain of ICT. Some stakeholders put forward the idea that robots may enjoy rights, while others suggest that they should have a specific legal status or even personhood.

Dual Use

Drones and robots that are used for surveillance, reconnaissance, and intelligence gathering are equipped with sensors, including cameras, infrared sensors, or lasers. They act as remote eyes for the observer: a better method of collecting information further afield and for longer periods of time. They reach out to areas that are difficult to access or potentially dangerous. The aim of data collection is to better anticipate failures, abnormal behavior, accidents, or attacks and to better implement a proper reaction. Such examples are the drones used for inspecting construction projects (dams, bridges, monuments, high-voltage power lines), monitoring high seas (pollution, pirates) or crowds (at sports events or demonstrations), gathering military intelligence, and distracting attention

during hostage-taking. Additionally, military robots may be equipped with weapons.

Presently, such robots are remotely operated or supervised by operators. It is, nevertheless, very likely that actual decision making will be shared between the robot and the operator or may even be delegated entirely to the robot. Researchers must focus on the increased capacities for autonomous situation recognition and autonomous decision making and on the associated risks, including perception errors; poor uncertainty assessment; and the difficulties in programming a common-sense, contextual, or moral judgment. One must deal with the problem of opacity of the robot's decisions and actions for the operator.

Based on its analysis of the applications of robotics and their impact on society, CERN has made several recommendations that we overview in the next section.

CERNA Recommendations

Autonomy and Decision-Making Capabilities

Context

Robot autonomy is the capacity to operate independently from a human operator or from another machine by exhibiting nontrivial behavior in a complex and changing environment. Programs governing the behavior of autonomous robots are designed to interpret sensory information, use this interpretation and prior knowledge to determine relevant actions, and compute the time and resources necessary for carrying out such actions. A distinction is to be made between robots that are supervised by an operator, i.e., a professional who possesses some knowledge of robot operation and is involved in decision making, and robots that interact with a user, i.e., a person with no special knowledge about robot implementation. With regard to this distinction, researchers should address robot autonomy in the context of the human-robot system rather than considering robots as isolated artifacts.

When a robot is supervised by an operator, its degree of autonomy belongs to a continuum spanning from complete human control to situations in which most functions are delegated to the machine while the operator only maintains high-level supervision or oversight [12]. In intermediate cases, some functions may require a human in the loop (e.g., interpreting images taken by a drone camera), while others, at the same time, can be delegated to the machine (e.g., drone navigation). Authority sharing is an important problem: researchers must decide whether a human or a robot holds the decision-making power at a given time or with regard to a particular function. Questions of control are crucial: When and how should the operator take over control of the actions previously commanded by the machine? Reciprocally, when and how should the machine take over control of actions commanded by the operator? Under which circumstances should the machine prevent the operator from retaking control (e.g., if the operator is impaired by workload, stress, or

emotions) [13], [14]? Are there objective and quantifiable criteria for such a takeover?

If all the decisions of a robot can possibly be delegated to the machine (e.g., when a quick reaction is necessary or in the absence of communication with the operator), researchers must address the issue of relevance and reliability of the knowledge and algorithms underlying the machine-made decisions and their limits. If an operatorless robot interacts with a user, particularly within the private sphere, its autonomy translates into a set of robotic functions that benefit the user. Researchers should address a crucial possibility for the user to disengage some functions or turn off the robot entirely. Under which circumstances is that possible? Furthermore, should the robot be able to prevent functional termination in a particular situation that it has evaluated using objective criteria? More generally, a human–robot system, however complex it may be, must be predictable and robust. In particular, researchers should address the problem of a trust bias in robots [15], user awareness [16], and the capacity of the whole system to deal with failures. Robot behavior must be traced to analyze malfunctions and report liabilities in the event of damage, injury, or loss.

Recommendations

Control

Researchers should investigate the capacity of the operator or the user to take control from the robot and that of the machine to take control from the human and specify the circumstances when such a takeover is allowed or mandatory. Researchers should investigate whether the human is to be allowed to disengage autonomous robotic functions.

Decisions Made Without the Operator's Awareness

Researchers must ensure that robotic decisions are not made without the operator's knowledge so as to avoid gaps in the operator's situational awareness. It is imperative to ensure that the operator will never believe the robot to be in a certain state while, in fact, it is in a different state.

Effects on the Operator's Behavior

Researchers should be aware of the trust bias, i.e., the operator's tendency to exhibit excessive confidence in robotic decision-making procedures, and of the moral buffer, i.e., the operator's tendency to morally disengage from robotic actions or behavior.

Programming Limits

Researchers should evaluate perception, interpretation, and decision-making software and make limitations explicit, i.e., the extent to which the models faithfully represent reality, the assumptions used in these models, and the criteria for computing a decision. Whenever a robot is endowed with moral behavior, researchers should evaluate whether general rules are applicable, if the notion of the right action is relevant to the moral framework used in

computation, and how moral values are ranked in controversial decision making.

Situational Awareness

With regard to interpretative robotic software, researchers should evaluate the extent to which the software can correctly characterize a situation and distinguish between apparently similar situations, in particular, in the circumstances when this characterization is the only basis of the ensuing decision or action.

Predictability of a Human–Robot System

Researchers should analyze the predictability of a human–robot system by considering uncertainty in any interpretation and action, possible robotic or human failures, and the entire set of states that can be reached by the system.

Traceability and Accounting

Researchers should develop tracing tools at the robot's design stage. These tools should facilitate accounting and the explanation of robotic behavior, even if this is done only in a limited way, at various levels intended for experts, operators, and users.

Imitation of Life, Affective, and Social Interaction

Context: Imitation of Living Beings

Biomimetic approaches consist in imitating living beings to better understand them or to obtain practical knowledge that will be put to use in artifacts. In the first case, biomimeticism increases knowledge by comparing biological reality with technical devices that reproduce some biological aspects. In the second case, the approach suggests new engineering solutions that seek to achieve the efficiency observed in nature. Certain research programs in robotics, e.g., microdrones with a flight scheme inspired by the flight of insects, use biomimeticism to a scientific or technological goal. A robot equipped with legs rather than wheels is better adapted for a human environment, such as a building with stairs.

Ethically speaking, the imitation of a living being implies a comparison between the artifact and the prototype. It serves to understand their degree of resemblance, to judge robot performance, and to address the possibly unintentional transposition of some of the prototype's features to the robot. To continue the previous example, in the case of a robot having legs rather than wheels, the projection of human or animal features is straightforward. It may blur the frontier between nature and artifact along the lines of what can be seen in film and fiction. Researchers should investigate the utility and necessity of the resemblance with living beings and make sure that they publicly address this issue in a clear manner.

Context: Emotion, Affectivity, and Attachment

Humans may respond emotionally as they interact with a robot. Although emotion is a hotly contested concept

[17], three modalities of the human emotional response remain uncontroversial. They include emotional expression, bodily symptoms and arousal, and subjective experience. *Affectivity* is the capacity to feel an emotion. *Affective computing* is the development of systems and devices that can recognize, interpret, process, and simulate human affects [18]. A machine that is able to interpret human emotional states adapts its behavior and provides an appropriate response. Robotic affectivity covers three technological aspects: a capacity to simulate what, in humans, corresponds to emotions (imprecisely referred to as *expressing emotions*); a capacity to understand human emotional expression; and a capacity to rationally take into account information that is contained in emotions. An affective robot that possesses these three functions interacts in different ways with different individuals.

The physical resemblance between a robot and a living being contributes to an emotional reaction, but anthropomorphization is not unusual, even for inanimate things. Other channels of resemblance include schematized mimicry, intonation of a sound or voice, or a playful modification of human/animal likeness, e.g., in a toy. Humans, too, sometimes project affectivity on nonaffective robots, e.g., on mine-clearing or vacuum-cleaning autonomous devices. Assistive robotics seeks to engage the user in an interaction with an affective robot to provoke the feeling of pleasure through such interaction and to enhance a level of trust in the robot. *Attachment* is an affective tie that is produced by attention given to another person or object, forging intimacy between the individual and the subject of attention. Only a handful of experiments have been conducted on the long-term use of affective companions and assistive robots. Researchers must consider the consequences of attachment, e.g., whether human dependence on a machine may become detrimental for contact with people.

Recommendations

Utility and Necessity in View of Purpose

Researchers should study the relevance and necessity of provoking emotions and of exhibiting biomimetic behavior or appearance in a robot, in particular, in the case of a strong visual or behavioral resemblance between a robot and a living being. When a human voice or likeness is imitated, researchers should investigate the effects of such imitation, including those exceeding the sphere for which the robot is intentionally designed.

The Nature–Artifact Frontier

Researchers should remain aware that a biomimetic approach may blur the frontier between nature and artifact. If a robotics project seeks almost perfect resemblance in any communication or perception channel between a robot and a living being, researchers must consult the operational ethics committee of their institution.

Study of the Effects

In an affective robotics project, researchers should investigate all consequences (e.g., social isolation or social stimulation) that their work may have on the user's ability to socialize with humans.

Child–Robot Interaction

In a robotics project that puts children in the presence of a robot, researchers should address the impact of child–robot interactions on the development of the child's emotional capabilities, particularly in early childhood.

Evaluation

In a robotics project that may involve user affectivity, e.g., by provoking attachment to a robot, researchers should draw up design and evaluation protocols and join with potential users and stakeholders in an effort to make the best informed scientific and technological choices.

Communication

Researchers should exercise caution when they speak in public on robot emotions and on their resemblance with living beings. Researchers should remain aware that emotional expression by a robot is an illusion in the human sense and that, intentionally or otherwise, the imitation of living beings may facilitate the transfer of certain features from the living being to the artifact.

Robot-Aided Therapy and Human–Robotic Augmentation

Context: Medical Robotics

Medicine and, in particular, surgery are major fields of application for robotics. Medical and surgical robots, artificial limbs and organs, and automatic systems for biological regulation are considered by state agencies as medical devices and are assessed accordingly through cost-benefit studies. General rules of medical ethics apply.

Repair devices reproduce, at least partially, functions that are absent in patients, such as the motor function for artificial legs or the prehensile function for artificial hands. In so doing, the devices increase the patient's independence and often allow the patient to survive or to maintain integrity. However, researchers must guarantee that such devices do not release personal information or cause vulnerability.

Context: Repair and Enhancement

Repairing humans, e.g., via bionic prostheses, is a matter of medical ethics, but, intentionally or otherwise, this technology can also generate enhanced capabilities or degrade natural human performance. Devices that play a reparative or a palliative role may be used, with small or no modifications, to enhance human functions by providing new capabilities that humans do not possess naturally. Exoskeletons, for example, are commonly used to

assist tetraplegics, but they can also enhance the bodily capacities of manual workers or soldiers.

Ethical opinions concerning human enhancement were published by the European Group on Ethics [19], [20] and by the Science and Technology Options Assessment Panel of the European Parliament [21]. A study by the National Science Foundation [22] lists societal issues and advocates public debate and deliberation. CERN recommendations are based on the principle that preserving autonomy, integrity, and independence, both for a repaired and for an initially healthy individual, must remain the researcher's imperative. Ethical issues include a possible generalized use of robotic augmentation techniques by healthy individuals wishing to enhance their performance in everyday life and the reversibility of such uses. If generalized, they will lead to a new societal situation.

Recommendations

Medical Ethics

Researchers in reparative or assistive robotics should, in coordination with health care professionals and patients, apply the principles of medical ethics to make informed choices between the requirements of care efficacy and safety, the independence and integrity of the patient, and privacy protection. These questions should be considered not only from the legal standpoint; ethical thinking and deliberation help to make individual adjustments on a case-by-case basis rather than apply a general rule. Researchers should solicit and follow opinions published by operational medical ethical committees.

Individual Independence and Integrity

Researchers working on reparative robotic systems should seek to preserve the independence of equipped individuals by situating them in a position to control their actions as extensively as possible. Researchers should also seek to preserve the integrity of other functions apart from those being repaired.

Reversibility

Researchers intentionally working on robotic devices for human enhancement must ensure that the resulting augmentation remains reversible. Devices should be removable without causing lasting harm or the loss of initial functions of the human body.

Societal Effects of Enhancement

Researchers should investigate the societal effects of human enhancement induced by the devices they develop, including the effects on the social behavior of equipped individuals and, reciprocally, on the social behavior of the unequipped.

Conclusion

When a robot leaves the laboratory to interact with people in a social context, it ceases to be a mere physical object and becomes a sociotechnical system. Thus, robot

ethics must address both the scientific and the societal aspects. On the societal side, robotics differs from other ICT areas in that it designs machines that are frequently used to represent the hopes and fears of humankind. Oftentimes these hopes and fears, amplified by media and science fiction, resound in a vague or excessive manner. On the scientific side, roboticists remain aware of the promises but also the limits of their discipline. An autonomous robot, for example, cannot be equipped with completely adequate ethical rules, and realizing perfect tracing or obtaining a full understanding of its behavior is an unreachable goal. An ethical attitude consists in connecting these two aspects of robotics in a coherent manner, through a deliberation involving other disciplines, in particular, human and social science, as well as the stakeholders and the users.

CERN recommendations indicate a road to be taken to enrich research with ethical reflection. They do not constitute a complete set of principles that every researcher should implement and do not impinge on the freedom of research. The role of research institutions is decisive for encouraging and supporting this approach in robotics and, more broadly, in information science and technology.

ACKNOWLEDGMENTS

We thank all CERN members and auditioned persons for their contribution.

REFERENCES

- [1] N. Wiener, *The Human Use of Human Beings: Cybernetics and Society*. Boston, MA: Houghton Mifflin, 1950.
- [2] N. Bostrom and M. M. Ćirković, Eds., *Global Catastrophic Risks*. London, U.K.: Oxford Univ. Press, 2008.
- [3] N. Bostrom, *Superintelligence: Paths, Dangers, Strategies*. London, U.K.: Oxford Univ. Press, 2014.
- [4] EURON. (2006, July). Roboethics Roadmap. [Online]. Available: <http://www.roboethics.org/atelier2006/docs/ROBOETHICS%20ROADMAP%20Rel2.1.1.pdf>
- [5] *IEEE Robot. Autom. Mag. (Special issue on roboethics)*, vol. 18, no. 1, Mar. 2011.
- [6] ETHICBOTS. (2007). ETHICBOTS home page. [Online]. Available: <http://ethicbots.na.infn.it/>
- [7] EU Project RoboLaw. (2014, Sept. 22). Final report D6.2: Guidelines on regulating robotics. Collaborative EC Project RoboLaw, Pisa, Italy. [Online]. Available: http://www.robolaw.eu/RoboLaw_files/documents/robolaw_d6.2_guidelinesregulatingrobotics_20140922.pdf
- [8] IEEE Standards Association. (2016). The IEEE Global Initiative for Ethical Considerations in Artificial Intelligence and Autonomous Systems. [Online]. Available: https://standards.ieee.org/develop/indconn/ec/autonomous_systems.html
- [9] R. Arkin, *Governing Lethal Behavior*. Boca Raton, FL: CRC Press, 2009.
- [10] European Parliament Committees. (2016). Committee on Legal Affairs. Working Group on Robotics and Artificial Intelligence. [Online]. Available: <http://www.europarl.europa.eu/committees/en/juri/subject-files.html?id=20150504CDT00301>

- [11] CERNA Opinion. (2014). Éthique de la recherche en robotique. (In French). Allistene, Paris, France. [Online]. Available: http://cerna-ethics-allistene.org/digitalAssets/38/38704_Avis_robotique_livret.pdf
 - [12] Department of Defense, Defense Science Board. (2012, July). The role of autonomy in DoD systems. Office of the Secretary of Defense, Washington, D.C. [Online]. Available: <http://www.acq.osd.mil/dsb/reports/AutonomyReport.pdf>
 - [13] N. Régis, F. Dehais, E. Rachelson, C. Thooris, S. Pizziol, M. Causse, and C. Tessier, "Formal detection of attentional tunneling in human operator—automation interactions," *IEEE Trans. Human-Mach. Syst.*, vol. 44, no. 3, pp. 326–336, 2014.
 - [14] S. Pizziol, C. Tessier, and F. Dehais, "Petri net-based modelling of human—automation conflicts in aviation," *Ergonom.*, vol. 57, no. 3, pp. 319–331, 2014.
 - [15] M. L. Cummings, "Automation and accountability in decision support system interface design," *J. Technol. Stud.*, vol. 32, pp. 23–31, 2006.
 - [16] N. D. Sarter, D. D. Woods, and C. E. Billings, "Automation surprises," in *Handbook of Human Factors and Ergonomics*, 2nd ed., New York: Wiley, 1997.
 - [17] L. Devillers, M. Tahon, M. A. Sehili, and A. Delaborde, "Inference of human beings' emotional states from speech in human–robot interactions," *Int. J. Social Robotics*, vol. 7, no. 4, pp. 451–463, 2015.
 - [18] L. Devillers, L. Vidrascu, and L. Lamel, "Emotion detection in real-life spoken dialogs recorded in call center," *J. Neural Networks*, vol. 18, no. 4, pp. 407–422, 2005.
 - [19] European Group on Ethics. (2005, Mar. 17). Ethical aspects of ICT implants in the human body, Opinion No. 20. [Online]. Available: http://ec.europa.eu/archives/bepa/european-group-ethics/docs/avis20_en.pdf
 - [20] European Group on Ethics. (2007). Ethical aspects of nanomedicine, Opinion No. 21. [Online]. Available: http://ec.europa.eu/archives/bepa/european-group-ethics/docs/publications/opinion_21_nano_en.pdf
 - [21] European Parliament STOA. (2015, May). Human enhancement study. Science and Technology Options Assessment Panel of the European Parliament, Strasbourg, France. [Online]. Available: https://www.itsa.kit.edu/downloads/etag_coua09a.pdf
 - [22] K. Scherer, T. Bänziger, and E. Roesch, *Blueprint for Affective Computing*. London, U.K.: Oxford Univ. Press, 2010.
- Alexei Grinbaum**, Alternative Energies and Atomic Energy Commission, Saclay, France. E-mail: alexei.grinbaum@cea.fr.
- Raja Chatila**, National Center for Scientific Research and University Paris-6 (University Pierre and Marie Curie), France. E-mail: raja.chatila@isir.upmc.fr.
- Laurence Devillers**, National Center for Scientific Research and University Paris-4 (Sorbonne), France. E-mail: devil@limsi.fr.
- Jean-Gabriel Ganascia**, National Center for Scientific Research and University Paris-6 (University Pierre and Marie Curie), France. E-mail: jean-gabriel.ganascia@lip6.fr.
- Catherine Tessier**, ONERA, Toulouse, France. E-mail: catherine.tessier@onera.fr.
- Max Dauchet**, National Center for Scientific Research, INRIA, and University of Lille, France. E-mail: max.dauchet@univ-lille1.fr.