On the Cusp: Computing Thrills and Perils and Professional Awakening

Natasa Milic-Frayling
Qatar Computing Research Institute
University of Nottingham
Intact Digital
natasa-milicf@frayling.net

ABSTRACT

Over the past eight decades, computer science has advanced as a field, and the computing profession has matured by establishing professional codes of conduct, fostering best practices, and establishing industry standards to support the proliferation of technologies and services. Research and applications of digital computation continue to change all aspects of human endeavor through new waves of innovation. While it is clear that different research advances fuel innovation, the ways they come together to make an impact vary. In contrast to highly regulated sectors such as pharma, medicine and law, the process of transforming research into widely deployed technologies is not regulated. We reflect on collective practices, from discovery by scientists and engineers to market delivery by entrepreneurs, industry leaders, and practitioners. We consider ecosystem changes that are required to sustain the transformational effects of new technologies and enable new practices to take root. Every such transformation ruptures in the existing socio-technical fabric and requires a concerted effort to remedy this through effective policies and regulations. Computing experts are involved in all phases and must match the transformational power of their innovation with the highest standard of professional conduct. We highlight the principles of responsible innovation and discuss three waves of digital innovation. We use wide and uncontrolled generative AI deployments to illustrate risks from the implosion of digital media due to contamination of digital records, removal of human agency, and risk to an individual's personhood.

PVLDB Reference Format:
Natasa Milic-Frayling. On the Cusp: Computing Thrills and Perils and Professional Awakening. PVLDB, 16(12): 4152 - 4159, 2023.
doi:10.14778/3611540.3611640

1 INTRODUCTION

Computing underpins the digital revolution and profoundly affects modern civilization. This is a bird's-eye view of innovation waves from the introduction of main-frame computing and first PCs to the internet, web, mobile technology, high-performance computing, the cloud, and AI. But how does that look on the ground? We reflect on our collective practices in the field of computing, from early innovation by scientists and engineers to market delivery by entrepreneurs, industry leaders, and practitioners and highlight the need to fulfil our professional mandate with full awareness of our responsibilities to ensure the safety, security, and protection of all those who are affected by our work.

Why is this important now? Over recent decades, we have matured as a profession, developed research and engineering practices and created technological wonders of the 20th and 21st centuries. We have enabled data capture and analyses that empower every science and industry, the delivery of information that serves billions of users in a split second, and the experience of real and virtual worlds that transform perceptions and livelihoods. The thrills of breakthroughs, insights and impact continue to drive our field, but the transformational power of our innovation must be matched by the highest standard of professional conduct.

We are intimately linked to the Information and Communication Technology (ICT) industry that leverages our research, innovation, and expertise to deliver capabilities of potential use across industries. In contrast to the controlled adoption of technology in regulated sectors such as pharma, health, manufacturing, and finance, with careful risk and quality management, ICT does not follow the same practices. Many ICT services aim at general consumption which often lacks a solid regulatory footing. In many instances, they do not apply principles of responsible deployment, including quality assurance and quality control for intended use. The perils of sidestepping our professional principles are amplified by the scale and impact of our technologies.

After 80 years of rapid innovation, we are on the cusp of a phase transition that can wipe out perhaps the most profound contribution to humanity: the use of the digital medium. We adopted it to capture, manage, and share human knowledge and made great efforts to protect the quality, truthfulness, and trust in its sources. The latest thrills of generative AI (genAI) carry great perils for digital humanity: uncontrolled use can turn knowledge records into undiscernible mush, destroy human agency in the digital world, and endanger human personhood in everyday life. This should be a turning point in our professional awakening: our professional ethics must be our guide and our professional code of conduct must be our utmost obligation. In the following sections we reflect on the computing research and innovation, professional ethics, and professional responsibility and shed light on a few key issues:

(1) Applications of technology by individuals and organizations induce accountability, moral responsibility, and liability. Technology itself does not.
(2) Without regulations, professionals and professional communities rely on professional code of conduct and ethics

Proceedings of the VLDB Endowment, Vol. 16, No. 12 ISSN 2150-8097.
doi:10.14778/3611540.3611640

to put in place policies and practices that adhere to ethical values and ensure public good.

(3) Sustainability and economic models are key to achieving technology impact through deployment, adoption and ongoing use. However, commercial professionalism is fragmenting our professional culture, affecting professional values, and disregarding public good.

(4) Quality assurance and quality control are essential aspects of deploying and enabling use of technologies that affect human lives. If the profession does not self-regulate it will lose autonomy implied by the social contract.

(5) Uncontrolled deployment of technologies can lead to irreversible damage within the digital ecosystem. Such is genAI pollution of digital information and the inability to ascertain entities and personhood within digital media, making the computing professional morally responsible for affecting the essential public good.

2 COMPUTING RESEARCH & INNOVATION

Research, as a pursuit of knowledge and discovery is fundamental to innovation. It provides necessary assets, insights, and know-how for inventions and practices that result in value for society. Computer science research has given rise to innovation through deployable applications, systems and services with transformative effects within specific sectors, from e-commerce, finance, and medical care to online entertainment and social media. It has also impacted the global ecosystem through infrastructures for computation and communication that made the digital medium and the digitalization of processes preferred over the analog.

2.1 Management of Research

As in other fields, computing research requires commitment to long-term investment and acceptance of uncertain outcomes and possibly intractable pathways of impact. The latter is particularly difficult in the computing field where innovation is realized through complex technical inventions that depend on research advances across different focus areas.

Traceability of impact. Within the academic sphere, there are established ways of sharing knowledge and insights through publishing and using citations to attribute work and recognize research contributions over time. Shared history is important for professional identity. Funding from national government and international funding agencies stratifies research grants into different categories, from fundamental to applied, from proof of feasibility and proof of concept to early prototype development and deployment. The principle is to create capacity for innovation to take place through semi-directed or undirected research.

In industrial settings, where publishing is not prioritized, research outcomes may never reach broader audiences but may support existing or new products and services. Instead of peer evaluation and approval, the utility of research is linked to product distribution and adoption and thus direct impact. While such directed research may be seen as a plausible way to drive innovation with a lower risk to investment, it bounds the areas of investigation and limits the scope of assets and know-how in the ecosystem. That would be hardly sufficient to enable breakthrough and transformational innovation.

Capacity for breakthroughs. The field of computing has evolved rapidly, with its foundations in electrical engineering and principles of digital computation. Researchers and engineers tackled diverse questions and application areas. Professional associations such as Association of Computing Machinery (ACM) and Institute of Electrical and Electronics Engineering (IEEE) have successfully enabled exchange of knowledge within specialty fields and cross-pollination of ideas by researchers involved in cross-disciplinary and multi-disciplinary initiatives. In order to grasp the scope of computing, it is instructive to look at the topic areas that are used by professional associations to keep track and provide access to works in the computing area. For example, IEEE Taxonomy (July 2023 v 1.02) for classification of research and engineering topics contains almost 7000 entries organized in three-level hierarchy of term-families. The taxonomy is a subset of a thesaurus that comprises 11,570 descriptive engineering, technical and scientific terms that represent concepts or units of thought presented in IEEE journals, conference papers, and industry standards [23, 24]

Novelty vs steadiness. Research publications as well as research investment typically require novelty and uniqueness which could possibly jeopardize much needed exploration of adjacent research topics or work that may be deemed incremental. However, the research community and its funders have been successful in justifying such efforts and preparing the ground for further stages of innovation.

2.2 Productization and Commercialization

Generally, technical aspects of computing innovation involve a combination of scientific principles, empirical evaluation, and engineering practices. Multiple assets and methods are pulled together to create proofs of feasibility and proofs of concept. Prototype systems may then be designed, motivated by envisioning the future or addressing specific issues that present challenges and opportunities for intervention. In either case, applied research remains locked within the academic realm, with limited impact, unless the application is a good market fit with a clear economic sustainability model. At a point where the technology is adopted and used, value is created. Some of that value may be captured as revenue and used to secure resources for sustained development, deployment and customer services. Once a steady revenue cycle is enabled, the technical innovation is on its way to affect a change and create new practices.

The transformational process from research to impactful innovation is complex, integrative and transdisciplinary—successful ventures necessarily involve experts from different fields and require diverse skill sets. This extended set of activities often lead to a perception of diminished contribution from research, difficulties with setting common objectives, and challenges in creating a productive work environment that accommodates different values, attitudes and practices. It is not surprising that for successful productization and commercialization, organizational leadership is key. It is also not surprising that efforts with diffused leadership and single-domain expertise, cannot progress beyond specific stages in the innovation process.

The EU government has, for example, put a great effort into promoting the use of open source technologies and solutions [5]. The open source community has provided a wealth of assets and created opportunities for system integrators to create affordable solutions. However, open source projects and the communities behind them lack mechanisms to capture the economic value of their work since they do not cover requirements beyond technical innovation. Sometimes a commitment to maintaining a system providing customer support services would be sufficient to retain users and translate the value of a deployed system into monetary rewards. Such income would increase the capacity for further production, maintenance, and customer support. However, organizing and providing legally binding support services is neither in the spirit of nor part of the open source project mandate.

Productization and commercialization. In an effort to guide progression from research to innovation, research administrators and practitioners have adopted various frameworks. Widely used are adaptations of NASA Technology Readiness Levels [2] which outline a journey from early research explorations to sustainable use of technical solutions. Such a framework is applicable to productization of research for use among specific stakeholders where the key objective is fit-for-use and continual investment from the stakeholders.

Bringing innovation to market requires a multi-faceted approach that covers both the development of technology, from prototype to product, and engagement with the ecosystem to ensure that the technology and practices take hold. The latter is typically out of scope of research and productization and involves expertise in finance, marketing, and sales. Over the years, we have seen several tactical approaches that seem amenable to developing tech enterprises, e.g., The Lean Startup [30], and tools to support specific activities, e.g., business canvas for conceptualizing and exploring a business opportunity and value canvas for sharpening the value proposition. Some of high-tech entrepreneurship wisdom has stood the test of time [14] and remains relevant. The work by Geoffrey Moore [18] on Crossing the Chasm points to a transition in the commercialization phases that pose high risk. More detailed analyses of the commercialization pathways are provided by the Triple Chasm model [28] which characterizes 3 important commercialization transitions and provides a detailed framework for assessing the stages and preparing for transitions. The Triple Chasm framework is particularly useful for researchers and engineers to relate technology readiness levels to commercialization readiness levels and assist with allocating efforts and resources appropriately to increase the chances of success. As professionals who understand the technology in depth, they have the privilege and responsibility to inform the design, implementation and deployment of products that are both fit-for-purpose and fit-to-market.

3 PROFESSION & PROFESSIONAL ETHICS

A common understanding of a profession is a vocation that requires education and practical experience in the field. Davis [8] elaborates and focuses on those who practice it: *"A profession is a number of individuals in the same occupation, voluntarily organized to earn a living by openly serving a certain moral ideal in a morally permissible way beyond what law, market, and morality would otherwise require."* Becoming a professional involves a process that, according to Ford & Gibbs [10], is well defined for fully developed professions. Such professions involve professional education, accreditation, skills development, certification, licensing, professional development, and a code of ethics. They provide a clear framework of engagement and a well-organized infrastructure to support existing members of the profession and to certify new ones. Key differentiators of mature professions are certification and licensing that determine who will be allowed to practice the profession.

Professional certification is a voluntary process: a non-governmental professional organization grants recognition to an individual who has met required qualifications. The certificate attests that the individual has demonstrated a certain level of mastery of a specific body of knowledge and skills within the relevant field of practice. On the other hand, licensure is a mandatory process where a government agency regulates a profession. The license grants permission to an individual to engage in an occupation if the applicant has attained the degree of competency required to ensure that public health, safety, and welfare will be reasonably protected. Once a licensing law has been passed it becomes illegal for anyone without a license to engage in that occupation.

In 1993 IEEE-CS and ACM set up a joint steering committee to explore the establishment of the software engineering profession. The committee conducted a survey of practitioners to understand the knowledge and skills required by software engineers, developed accreditation criteria for undergraduate programs in software engineering, develop a code of ethics for software engineers. Furthermore, in May 1999, ACM Council passed a resolution: "*ACM is opposed to the licensing of software engineers at this time because ACM believes that it is premature and would not be effective in addressing the problems of software quality and reliability.*"

Exceptions are engineers who work in the United States on systems that are deemed critical and high risk for the nation in which case an engineering license is required. The National Council of Examiners for Engineering and Surveying (NCEES) offers professional licensure in Electrical and Computer engineering, covering the Fundamentals of Engineering (FE) exam for recent graduates and students who are close to finishing an undergraduate engineering degree from an EAC/ABET-accredited program and the Principles and Practice of Engineering (PE) exam as a minimum level of required competency for engineers who have gained a minimum of four years post-college work experience in their engineering discipline [25, 26].

3.1 Professionalism & Social Contract

Considering that the ICT industry and electrical and computer engineering are not regulated, it is important to reflect on professionalism as a social contract between professionals and the public. Generally, society uses the concept of a profession to organize the delivery of essential services that it requires. Under the terms of implied "social contract", a profession is given autonomy and the privilege to self-regulate in return for being trustworthy, assuring the competency of its members, and dedicated to the public good. Every member of the profession becomes part of the "community of practice" during education and training, accepting the norms and values of the community and, in return, acquiring a professional identity that the community confers upon them. The leaders of the community continually negotiate the social contract on behalf of the profession and maintain its collective professional status. Upholding and communicating the importance of its devotion to the public good is critical for securing the profession's status.

This ideology of professionalism has been explored in depth by the legal and medical professions since besides the regulatory framework it is the public trust in the profession that is of utmost importance [7, 29]. In reality, competitive market practices are increasing the risk of eroding trust in professionalism. In the case of a crisis, the public will attempt to renegotiate the social contract and the government will step in, limiting the profession's self-regulation. At the individual level, market ideology affects the very core of the profession's cohesion, where competition among professional organizations takes precedence over both the professional collegiality and public welfare.

According to *structural functionalism*, society is a complex system that evolves to establish specialized parts, each with a specific function to achieve common goals [6]. Increased competition and a rise of *commercialized professionalism* points to a society that operates based on social conflict theory where groups attain differing amounts of material and non-material resources through various forms of conflict, and more powerful groups use their power to retain it and exploit groups with less power. In the case of commercial service, individual organizations strive to ensure "*ongoing reliance of the community on their services for their own self-preservation*" (Kelly 1994, p.267 [13]) while possibly affecting the resilience and stability that society needs. As a result, professional values cannot be defined any more with regard to the common societal good and independently of specific organizations, and the professional culture becomes fragmented into commercialized professional subcultures.

3.2 Code of Professional Conduct

Members of both ACM and IEEE have formulated the Code of Ethics and Code of Conduct [1, 21, 22]. They typically cover principles, rules, ideals, requirements, permissions, and prohibitions, with the aim of guiding members to act responsibly, reflect on the societal impacts of their work and make a concerted effort to support the public good. For example, the ACM Code of Ethics and Professional Conduct is aimed at all computing professionals and those who use computing technologies in impactful ways [1]. It includes principles that describe professional responsibilities and guidelines for ensuring the public good. It also outlines guidance for members in leadership roles in the workplace and as ACM volunteers.

The professional community recognizes the limitations of the Code for solving ethical problems that one may face. Professional conduct cannot be reduced to rule-based professionalism. Indeed, *"Codes of ethics suffer the same fundamental problem as ethical theories—goodness cannot be defined through a legalistic enumeration of do-s and don't-s; it must come from the heart."* (Bowyer 2001 [4]). However, the Code serves as a basis for ethical decision-making. For a given issue, one can identify one or more principles that should be taken into account and carefully consider how the situation relates to public good. By having a common framework for reasoning about ethical issues and engaging in open discussions, the computing profession benefits from transparency and accountability to all the stakeholders.

3.3 Common Values

Creating a productive working environment with experts from different professions requires a concerted effort to establish a common ground for professional conduct. It is a common practice to establish fundamental principles based on Virtue Ethics that stem from Aristotle:

- Be impartial
- Disclose information that others ought to know
- Respect the rights of others
- Treat others justly
- Take responsibility for your actions and inactions
- Take responsibility for the actions of those you supervise
- Maintain your integrity
- Continually improve your abilities
- Share your knowledge expertise and values.

Advantages of the Virtue Ethics are in its motivation for good behavior and healthy social interactions that are widely accepted and underpin professional ethics across domains.

4 RESPONSIBLE INNOVATION

Computing professionals strive to deliver dependable computing system. As part of their professional remit, they have moral responsibility to provide applications, systems, and services that can justifiably be trusted. With the public good in mind, they ought to minimize potential injury, danger, and catastrophic consequences to users and the environment [15, 16].

4.1 Professional Standards and Best Practices

The dependability of computing systems and services is characterised by

- Availability—Readiness of correct service
- Reliability—Continuity of correct service
- Safety—Absence of catastrophic consequences on the user(s) and the environment
- Confidentiality—Absence of unauthorized disclosure of information
- Integrity—Absence of improper system state alterations
- Maintainability—Ability to undergo repairs and modifications.

These are achieved by applying best engineering practices. Development of software systems, in particular, has been optimized over years through research of engineering practices that led to theories, methodologies, and tools for supporting software production in all its stages . The waterfall method, for example, comprises:

- Specification—Defining the functions to be performed by the software
- Development—Producing the software that meets the specifications
- Validation—Testing the software
- Evolution—Modifying the software to meet the changing needs of the users.

Considering system validation, arriving at an effective process for software testing is a challenge due to the complexity of the systems design and implementation. The aim is to assess the correctness, completeness, and quality of the developed computer software and ensure that it satisfies the specification and user needs. Testing can reveal bugs but cannot prove that the program will work correctly under all circumstances. Formally proving that the software meets specifications is costly and often infeasible. Even if one can prove that the program is 'correct' and meets specifications, there is uncertainty whether the specifications are correct. This fundamental challenge in software engineering is well summarized by Tony Hoare, Turing Award Winner 1980 [11]: *"There are two ways of constructing a software design. One way is to make it so simple that there are obviously no deficiencies. The other way is to make it so complicated that there are no obvious deficiencies."*

In practice, engineers apply software testing principles and skills they develop over time and continuously improve. Fundamentally, computer systems are only as good as humans who make them. As noted by Tony Hoare [12]: "The real value of tests is not that they detect bugs in the code, but that they detect inadequacies in the methods, concentration, and skills of those who design and produce the code." Computing system errors stem from human errors and are typically traceable. In order to uphold the highest professional standards, it is critical for the professional community to foster the culture of openness, knowledge sharing, and transparency about professional accountability. While accountability for erroneous system can be determined relatively easily, the moral and legal responsibilities are much harder to establish.

4.2 Accountability, Responsibility & Liability

Over the past decades, the computing profession has engaged with standardization bodies, developing ICT standards and frameworks that pertain to various aspects of computing. Software application development has been covered by Capability Maturity Model Integrated (CMM/CMMI); technical standards concerning interoperability, metadata, and web-related technologies by Organization for the Advancement of Structured Information Standards (OASIS), Object Management Group (OMG), W3C, and Dublin Core; overall ICT governance and management by COBIT®; principles of good corporate IT governance by ISO/IEC 38500; data management by DAMA-DMBOK, to name a few.

However, the professional community has neglected accountability for the impact of computing, specifically for the harms and risks of faulty and malfunctioning systems. This is, for example, reflected in the wide-spread practices in the ICT industry of providing software warranties that leave users with no recourse in the case of faults and errors. Since accountability is, in essence, answerability, i.e., a state of being compelled to or called to account for one's action, it applies to everyone who is involved in a specific action. Thus, even if things go terribly wrong with a computing system and users have no recourse on the ground of system warranty, they can be assured of answerability.

A stronger focus on accountability is important for the profession in order to increase global awareness, introspection, and self-regulation. It is the first step towards attaining a sense of responsibility as a virtue. In terms of social welfare and public good, a culture of accountability motivates actions to prevent harm and minimize risks. It also provides a reasonable starting point for issuing penalties and punishment, and securing compensation for victims of harm caused by failures.

In order to promote accountability, it is important to understand reasons for the current lack of its adoption. Research by Nissenbaum [20] identified attitudes that serve as "barriers" to the culture of accountability:

- Many hands—Computing systems are built by teams, complex, multi-layered, and dependent on other systems. Thus, in tracking the faults it is difficult to assign responsibility.
- Bugs—The view that bugs are inevitable implies that, while harms and inconvenience are regrettable, they cannot be helped, and it would be unreasonable to keep programmers responsible.
- Computer as Scapegoat— People point at the complexity of the computer to argue it was the computer's fault when things go wrong.
- Ownership without Liability—Commercial companies protect computing innovation (IP) and take advantage of exclusive use, without responsibility to protect from harm.

Accountability is inferred from the nature of an action and from the relationship of a person, i.e., the acting agent to the outcomes of that action. In many instances, accountability is arbitrated by investigating "causal" and "fault" conditions and determining blameworthiness.

Liability, on the other hand, is primarily focussed on a person or organization who is to blame and needs to compensate victims for damages suffered after an undesirable event. Liability is rooted in the suffering of victims and the starting point for assessing liability is the victim's condition. From the perspective of professional practices, focussing solely on liability and compensations to victims, is not conducive to improving professional practices and removing the cause of harm.

Moral responsibility of individuals or a group is considered when their voluntary actions have morally significant outcomes which would make it appropriate to blame or praise them. Ascribing moral responsibility establishes a link between the subject, i.e., a person or a group of people, and the object, i.e., someone or something, that is affected by the actions of the subject. This can be done retrospectively and prospectively and due to the complexity of computing systems and stakeholder involvements, it raises several important questions: when is it appropriate to ascribe moral responsibility, what awareness and freedom of will is required, and are humans the only entities to which moral responsibility can be attributed? It is accepted that

- There should be a causal connection between the person, i.e., the subject, and the outcome of actions
- The subject has to have knowledge of and be able to consider the possible consequences of performed actions
- The subject has to be able to freely choose to act in a certain way.

Birsch [3] offers criteria for moral responsibility: For a person to be morally responsible for the consequences of computing system failure:

(1) The action of the person must have caused the harm or have been a significant causal factor.
(2) The person must have intended or willed the harm, or it must be a result of his or her negligence, carelessness or recklessness.
(3) The person must have been able to have known, or must know of the consequences of the action, or must have deliberately remained ignorant of them.

Computing practitioners have a tendency to ignore, side-step, and avoid responsibility, often due to misconceptions about the nature of their work and the notion of moral responsibility [19]. Two pervasive ones are that

- Computing is an ethically neutral practice. This stems from a narrow, technology-centred focus on the development of computing systems and from ignoring the broader context of technology use.
- Responsibility is only about determining blame when something goes wrong.

Unfortunately, focusing on the malpractice model of responsibility is likely to deepen the avoidance of responsibility. Since developers' work is often distant from technology deployment and use, that provides a basis for claiming that there is no direct relationship and thus a diminished causal link to any malfunction. Instead, it is better to focus on positive responsibility, emphasizing the virtue of having a deep regard for the consequences that actions may have on others. That shifts the attention towards minimizing foreseeable undesirable events and taking actions to prevent them rather than focusing on blaming or punishing irresponsible behavior.

4.3 Human Agency and Artifact Agency

With increased adoption of automation, there are concerns about the control of computing systems and human agency in developing and using such artifacts. According to Verbeek [31], all forms of human action can be related to three forms of agency:

- The agency of the human performing the action
- The agency of the designer who helped shape the mediating role of the artifacts
- The artifact mediating human action.

The agency of artifacts is inextricably linked to the agency of its designers and users, but cannot be reduced to either of them. A subject that acts or makes moral decision is a composite of human and technological components. Verbeek claims that moral agency is not merely located in a human being but in a complex blend of humans and technologies.

In the computing field, it is essential to set a normative guide for professionals who design, develop, deploy, evaluate or use computing artifacts. The term computing artefact refers to any artefact that includes an executing computer program. A moral responsibility for computing artefacts would imply that professionals who produce them or use them are answerable for their behavior and effects. However, that computing artifacts would need to be considered in the context of socio-technical systems comprising people, artefacts, physical surroundings, customs, relationships, assumptions, procedures and protocols.

In 2010, an interdisciplinary group of philosophers, computer scientists, practitioners, and lawyers approached these issues and began developing *"Principles Governing Moral Responsibility for Computing Artifacts"* (*"Rules"*) presented by Miller [17] as five rules:

Rule 1: The people who design, develop, or deploy a computing artifact are morally responsible for that artifact, and foreseeable effects of that artifact. This responsibility is shared with other people who design, develop, deploy or knowingly use the artifact as part of a sociotechnical system.

Rule 2: The shared responsibility of computing artifacts is not a zero-sum game. Responsibility is not reduced because more people become involved. A person's responsibility includes being answerable for behaviors and effects after deployment, to the degree to which these effects are reasonably foreseeable by that person.

Rule 3: People who knowingly use a computing artifact are morally responsible for that use.

*Rule 4:*People who knowingly design, develop, deploy, or use a computing artifact can do so responsibly only when they make a reasonable effort to take into account the sociotechnical systems in which the artifact is embedded.

Rule 5: People who design, develop, deploy, promote, or evaluate a computing artifact should not explicitly or implicitly deceive users about the artifact or its foreseeable effects, or about the sociotechnical systems in which the artifact is embedded.

5 HINDSIGHT & FORESIGHT

Similar to previous transformations of humanity, digital revolution has brought about changes that have deeply affected society. With every change there are advantages and pitfalls that society needs to deal with. Here we reflect on three waves of innovation with transformational effects on individuals and society and point to one common aspect. All three waves are driven by the ICT industry with computing professionals playing instrumental roles in bringing them about. However, computing professionals failed to moderate negative effects and prevent harm. Due to the lack of self-regulations, governments and international bodies had to intervene and impose polices and regulations.

5.1 Search, Advertising, and Privacy

With the emergence of the Internet and World Wide Web we have established a global communication and information publishing infrastructure that has grown to support large scale services in e-commerce, social networking, and social media. It is instructive to see how computing professionals from researchers to engineers, entrepreneurs and business leaders have engaged and carried their professional responsibility.

Early online search engines had to deal with the scale of crawling published material, suboptimal internet connections, socially unacceptable material, and volumes of user queries. Most critically, they struggled to identify effective business models that could sustain search operation and other services and deliver benefits from open information access to users. Internet browsers became a critical component within the distributed architecture in which content is independently managed on web-sites and full user context can only be captured within the browsers. A number of vendors attempted to provide added-value through toolbars that enhance the user experience by filtering content, or by providing access to specialized high-value content. Soon a business model emerged thanks to e-commerce and Web advertising. The simple invention of using cookies to maintain a session between client and Web server was used for the purposes of tracking users as they navigated the Web.

Advertising services created effective ad-bidding platforms that within milliseconds fulfilled ad-placement orders and displayed ads in the context of Web pages viewed by targeted users. Access to free information on the internet was paid for by users indirectly through their digital footprints that were captured automatically using browser cookies. Users were unaware of the value exchange due to the lack of transparency of the browser settings related to cookie management.

The issues of privacy violations became more apparent with the prevalence of targeting advertisements and expansion of the privacy and security attacks that utilized information from online browsing. Increases in identity thefts and financial scams became a serious problem. Potential for extensive surveillance across devices and services became easy to do thanks to the same technology that was put in place to enable easy use and convenience such as single sign-on. With governments themselves being hit by international scandals of leaked information and increased awareness of the privacy violations, they stepped in to regulate private data retention and usage. GDPR was established in the EU and instigated a number of new practices that limit, to a degree, user tracking on Web sites [27]. However, the user has to assume that, unless using browsing in a private window, they are fully exposed, and information of their online navigation will be used in a non-transparent way.

5.2 Social Media and Societal Changes

Emergence of social media platform represented another engineering marvel with the volumes of data and user engagements at unprecedented scales. The design of interfaces and interaction models encouraged networking and information propagation. The business model involved digital footprints of the users with little transparency of the value chain that users were inadvertently fueling based on their content and social interactions. From Facebook, Twitter, LinkedIn to Instagram, Snapchat and TikTok, the wave of social media services transformed communication and human interaction, affecting the social fabric in a positive and negative ways. Broad deployment and unrestricted access amplified all aspects of social interactions from the use of the platforms for political campaigns and political movements, often accompanied by propaganda and misinformation duels, to cyber-bullying and ghosting that resulted in loss of lives. Due to lack of self-regulation, governments again had to step in and demand protection of users, developing policies and legal frameworks to process afflicted harm.

5.3 AI, Human Agency and Personhood

Most recent advances in deep learning models have begun the next wave of innovation and digital revolution. Machine Learning and AI methods have been used to address needs for classification, prediction, optimization, and automation of systems across industries. AI techniques have been used to identify promising patterns in drug discovery, manufacturing design, and operations optimization. Deep learning techniques have been applied successfully to enable machine translation, speech recognition and voice synthesis. With controlled deployments within specific sectors, the technology is used within the boundaries of regulations and industry standards. However, recent unrestricted releases of proprietary and open-source large language models (LLMs) that fuel generative AI services, like ChatGPT, are raising new concerns. While the outputs of generative AI (genAI) services seem coherent, the systems cannot guarantee that the output is actually correct. In computing research, it is common to release technology probe to see how specific technology could be used [9]. The release of ChatGPT is akin to a technical probe since besides the general statement about the chat use, the system specification and intended use are not presented or documented. Even the interaction model through user prompts is a matter of discovery.

Due to the very nature of the technology, content authored and validated by humans, cannot be easily discerned from the output of a genAI system that is produced unchecked. This immediately raises concerns for the contamination of digital records and diminishing role of human agency in constructing and sharing information. Furthermore, advances in speech and vision technologies will make it possible to create avatars and synthetic voices of such a high quality that, in the digital realm, they will be indiscernible from a real person. The loss of personhood in the digital sphere will be detrimental for digital communication and may just be the last straw in the collapse of trustful online interaction.

6 CONCLUDING REMARKS

Researchers and engineers in computing are part of innovation waves and as computing professionals need to stay aware and cognizant of the implications of their activities for the public good. If working for organizations that are actively pursuing commercialization, it is important they find common ground with professionals from other fields to deliver effective and dependable technologies and solutions. Activities and decision of individuals are supported by professional code of conduct that outlines the principles of professional behavior. They must resist the trends of commercialized professionalism that puts the interests of individual organizations above social good and professional values. Computing professionals are accountable to the public for their actions and answerable for the impact and implication of technology use. Moral responsibility is not always easy to establish due to the complexity of processes and systems but there are frameworks that can be used effectively to guide and assess professional actions.

With increased adoption of automation and computing artifacts with imbedded computation, there are tendencies to transfer accountability and responsibility onto the artifacts. This needs to be resisted and, instead, a principled way of considering full socio-technical context and all relevant stakeholders must be developed. Finally, the computing profession is operating on the cusp of a qualitative change in the digital ecosystem that can dramatically affect the use of digital media, at least for information publishing and communication. Thus, adhering to the engineering principles that ensure quality assurance and quality control is of utmost importance. The wide distribution and uncontrolled use of technologies that can deliberately or unintentionally pollute digital records and affect digital personhood, can destroy the utility of the digital medium and wipe out the most important achievements of the digital revolution.

SPEAKER BIOGRAPHY

Natasa Milic-Frayling. Dr Natasa Milic-Frayling is Research Director at the Qatar Computing Research Institute (QCRI) leading research and innovation initiatives of the Arabic Language Technology group. She is Professor Emerita with the School of Computer Science at the University of Nottingham where she served as Chair of Data Science and taught professional ethics in computing. In 2016, Natasa founded Intact Digital Ltd, a digital continuity company that addresses the issues of digital obsolescence and long-term care of software technologies in order to protect valuable digital assets. As Intact Digital CEO, she is working with decision makers and thought leaders in the pharma sector and life-sciences to ensure data integrity and reconstruction of research studies for decades to come.

Natasa has 25+ years of experience in computer science research and innovation, including 17 years at Microsoft Research (MSR) where she led Integrated Systems group, focusing on interdisciplinary research in information management and communication. She has authored research publications across domains, from information retrieval, machine learning and dialogue generation to security and HCI, and has a dozen approved patents to her name. Besides her research role, Natasa led MSR Research Partnership Programme, promoting collaboration on strategic ICT industry challenges, including digital preservation.

Natasa is actively engaged with a broader professional community. She is a member of the Preservation Sub-Committee within the UNESCO Memory of the World Programme and serves as Chair of the Research and Technology Working group for the UNESCO PERSIST project. She is a member of the ACM Europe Technology Policy Committee and IT Committee of the Research Quality Association (RQA). In the past she has served on the ACM Europe Council and as Chair of the ACM Women Europe Executive Committee.

REFERENCES

[1] 2020. *ACM Code of Ethics and Professional Conduct.* https://www.acm.org/code-of-ethics

[2] National Aeronautics and Space Administration (NASA). 2012. *NASA Technology Readiness Level.* https://www.nasa.gov/directorates/heo/scan/engineering/technology/technology_readiness_level

[3] Douglas Birsch. 2004. Moral Responsibility for Harm Caused by Computer System Failures. *Ethics and Information Technology* 6, 4 (2004), 233–245. https://doi.org/10.1007/s10676-005-5609-5

[4] K. Bowyer. 2001. *Ethics and Computing: Living Responsibly in a Computerized World* (2nd. ed.). Wiley-IEEE Press.

[5] European Commission. 2020-2023. *EU Open Source Software Strategy.* https://commission.europa.eu/about-european-commission/departments-and-executive-agencies/informatics/open-source-software-strategy_en

[6] Lillian Corbin. 2005. How "Firm" are Lawyers' Perceptions of Professionalism? *Legal Ethics* 8, 2 (2005), 265–290. https://doi.org/10.1080/1460728X.2005.11424241

[7] Sylvia R. Cruess and Richard L. Cruess. 2004. Professionalism and medicine's social contract with society. *AMA Journal of Ethics* 6, 4 (2004), 185–188.

[8] M. Davis. 2002. *Profession, Code, and Ethics.* Ashgate. 256 pages.

[9] Hilary Hutchinson, Wendy Mackay, et al. 2003. Technology Probes: Inspiring Design for and with Families. *Conference on Human Factors in Computing Systems - Proceedings* (April 2003), 17–24.

[10] Gary Ford and Norman Gibbs. 1996. *A Mature Profession of Software Engineering.* Technical Report CMU/SEI-96-TR-004. Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA.

[11] Charles Antony Richard Hoare. 1981. The Emperor's Old Clothes. *Commun. ACM* 24, 2 (Feb. 1981), 75–83. https://doi.org/10.1145/358549.358561

[12] C. A. R. Hoare. 1996. How Did Software Get So Reliable Without Proof?. In *FME '96: Industrial Benefit and Advances in Formal Methods, Third International Symposium of Formal Methods Europe, Co-Sponsored by IFIP WG 14.3, Oxford, UK, March 18-22, 1996, Proceedings (Lecture Notes in Computer Science)*, Marie-Claude Gaudel and Jim Woodcock (Eds.), Vol. 1051. Springer, 1–17. https://doi.org/10.1007/3-540-60973-3_77

[13] M.J. Kelly. 1996. *Lives of Lawyers: Journeys in the Organizations of Practice.* University of Michigan Press.

[14] J. Lang and Cambridge Entrepreneurship Centre. 2002. *The High-tech Entrepreneur's Handbook: How to Start and Run a High-tech Company.* Pearson Education.

[15] Phillip A Laplante. 2004. First, Do No Harm: A Hippocratic Oath for Software Developers? What's Wrong with Taking Our Profession a Little More Seriously? *Queue* 2, 4 (Jun 2004), 14–18. https://doi.org/10.1145/1016978.1016991

[16] Steve McConnell and Leonard Tripp. 1999. Guest Editors' Introduction: Professional Software Engineering-Fact or Fiction? *IEEE Softw.* 16, 6 (nov 1999), 13–18. https://doi.org/10.1109/MS.1999.805468

[17] Keith Miller. 2011. Moral Responsibility for Computing Artifacts: The Rules. *IT Professional* 13 (07 2011), 57 – 59. https://doi.org/10.1109/MITP.2011.46

[18] G.A. Moore. 1995. *Crossing the Chasm: Marketing and Selling High-tech Products to Mainstream Customers.* HarperBusiness.

[19] George V. Neville-Neil. 2014. Outsourcing Responsibility. *Commun. ACM* 57, 10 (Sep 2014), 28–29. https://doi.org/10.1145/2661051

[20] Helen Nissenbaum. 1996. Accountability in a computerized society. *Science and engineering ethics* 2, 1 (1996), 25–42. https://doi.org/10.1007/BF02639315

[21] Institute of Electrical and Electronics Engineers (IEEE). 2014. *IEEE Code of Conduct.* https://www.ieee.org/content/dam/ieee-org/ieee/web/org/about/ieee_code_of_conduct.pdf

[22] Institute of Electrical and Electronics Engineers (IEEE). 2020. *IEEE Code of Ethics.* https://www.ieee.org/content/dam/ieee-org/ieee/web/org/about/corporate/ieee-code-of-ethics.pdf

[23] Institute of Electrical and Electronics Engineers (IEEE). 2023. *IEEE Taxonomy.* https://www.ieee.org/content/dam/ieee-org/ieee/web/org/pubs/ieee-taxonomy.pdf

[24] Institute of Electrical and Electronics Engineers (IEEE). 2023. *IEEE Thesaurus v.1.02.* https://www.ieee.org/content/dam/ieee-org/ieee/web/org/pubs/ieee-thesaurus.pdf

[25] National Council of Examiners for Engineering and Surveying (NCEES). 2022. *NCEES Fundamentals of Engineering (FE) Electrical and Computer CBT Exam Specifications.* https://ncees.org/wp-content/uploads/2022/09/FE-Electrical-and-Computer-CBT-specs.pdf

[26] National Council of Examiners for Engineering and Surveying (NCEES). 2023. *NCEES Principles and Practice of Engineering (PE) Electrical and Computer Exam.* https://ncees.org/exams/pe-exam/electrical-and-computer/

[27] The European Parliament and Council. 2016. *EU Data Protection Rules.* https://commission.europa.eu/law/law-topic/data-protection/eu-data-protection-rules_en

[28] Uday Phadke and Shailendra Vyakarnam. 2018. *Scale-up Manual, The: Handbook For Innovators, Entrepreneurs, Teams And Firms.* WORLD SCIENTIFIC (EUROPE). https://doi.org/10.1142/q0176

[29] Fred Phillips. 2016. *Ethics of the Legal Profession.* Routledge.

[30] E. Ries. 2011. *The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses.* Crown.

[31] Peter-Paul Verbeek. 2006. Materializing Morality: Design Ethics and Technological Mediation. *Science, Technology, & Human Values* 31, 3 (2006), 361–380. https://doi.org/10.1177/0162243905285847