Deep Science Publishing
https://doi.org/10.70593/978-81-989050-0-0_12

Chapter 12: Enabling the next generation of human-centered services through autonomous decision support

12.1. Introduction

A distinguishing feature of the emerging digital ecosystem is that value is created and exchanged, as never before, not just among businesses and consumers, but among other individuals as well as governments and civil society, who often act as the ultimate arbiters of transnational market relations. Digital platforms enable these interconnections. The digitization of tangible resources that were traditionally controlled by companies generates new assets for the platform economy, content generated by users and customers. The ever-increasing integration and mutual reinforcement of service delivery and consumption creates new opportunities for business, stronger exchanges among co-producers and consumers, and more engaging and personalized user journeys. Automating the increasing complexity of this process is essential to enable the next generation of human-centered services (Kamar, 2016; Endsley, 2017; Gajos et al., 2018).

The current status of Artificial Intelligence development makes it possible to construct adaptive solutions to the emerging complexities of digital services in these two critical areas. AI-nurtured services can easily accompany customers throughout the engagement lifecycle, from pre-purchase search and selection to the actual transaction, from payment to delivery, and post-purchase use, feedback, and complaint support. AI can also help automate service delivery to individual customers by providing customer-driven and -controlled push services, targeted micro-services, personalized offers and experiences, context-aware and genome-centered products and services, and requests for information and long-term relationship systems. In conclusion, adaptive solutions enabled by Artificial Intelligence can facilitate the next generation of personal and personalized human-centered services. They can help augment and automate traditional enterprise and interactional activities (Smith et al., 2009; Susskind & Susskind, 2015; Nguyen et al., 2022).

12.1.1. Overview of the Document's Focus and Scope

The focus of this document is on the integration of AI decision support services into highly interconnected human-problem systems. A particular emphasis is given to those systems that have been historically characterized by reliance on a human workforce, where centuries of pragmatic domain practice have led to exceptionally highly tuned and resilient perception and reasoning faculties. While that reliance will continue for the foreseeable future, it is gradually becoming apparent that a combination of humans and machines can achieve a performance that is better than either alone. Thinking in terms of a greater-than-the-sum-of-the-parts performance suggests that the autonomous and semi-autonomous decision-aid providers that will be created in the next few years should be consciously designed to enhance and exploit distinctively human faculties while alleviating the shortcomings that humans possess. These collaborative partnerships should yield mutual benefits, with less exhausting workloads and greater performance for the human workers, while permitting the service organizations to reduce costs, increase quality, and expand their services.

Humans are not "bottlenecks" to be removed from decision-making tasks, but rather experts within loops that need to be carefully designed to yield enhanced cognitive capabilities and reduced workload. However, despite our profound empathy with the hard decisions that many of these services are required to perform daily, the justification for developing the decision-aid systems that we advocate comes not from attempts to alleviate the psychological burden that some of these decisions entail, but rather from a thorough grounding in economic imperatives. It is the steadily growing cost of domain expertise and performance that is coming to characterize decision-centric services in the 21st-century economy that we view as being the driving force behind the practical impetus for AI and other autonomous, distinctive support services.

12.2. Understanding Human-Centered Services

1. Definition and Importance A key demographic trend for modern societies, particularly in the western hemisphere, is the aging of the population and the subsequent increase of the dependency ratio. This increase leads to more people needing services, while there are fewer and fewer active members of society to provide those services. In this context, Digital Services have become a frontier for research, but also for local or global companies providing digital solutions. Digital Services can be defined as services that are either fully or partially service digital technology-mediated. The explosion of openly available cloud and access-based digital services tools laden around artificial intelligence and machine learning have allowed the development of a plethora of digital services across domains. Health and wellness, information and knowledge, and learning and growth services are some examples of the usage of these technologies. Digital Services

are capable of providing the scaling that traditional services reliant on human activities alone are unable to deliver. These Digital Services can provide services unavailable to people in rural or excluded communities. For greater or lesser socioeconomic dynamics, they can further contribute to bridging the social and personal divide of loneliness. Yet, the availability of Human-Centered Digital Services does require the structural and infrastructural conditions through private and public investment.

2. Historical Context The Industrial Evolution changed the focus of services from a traditionally human-centered endeavor, for example in craftsman and artisanal products, to a more automated, product-oriented approach. Goods were mass-produced, and the focus shifted to increasingly large-scale production, often resulting in celebrity products. The Digital Revolution has seemingly reversed this trend. New technology has again enabled a more personalized focus on services. Instead of being perceived as a cost center, Human-Centered Services are now seen as being at the center as their quality is being used to differentiate products in increasingly competitive markets.

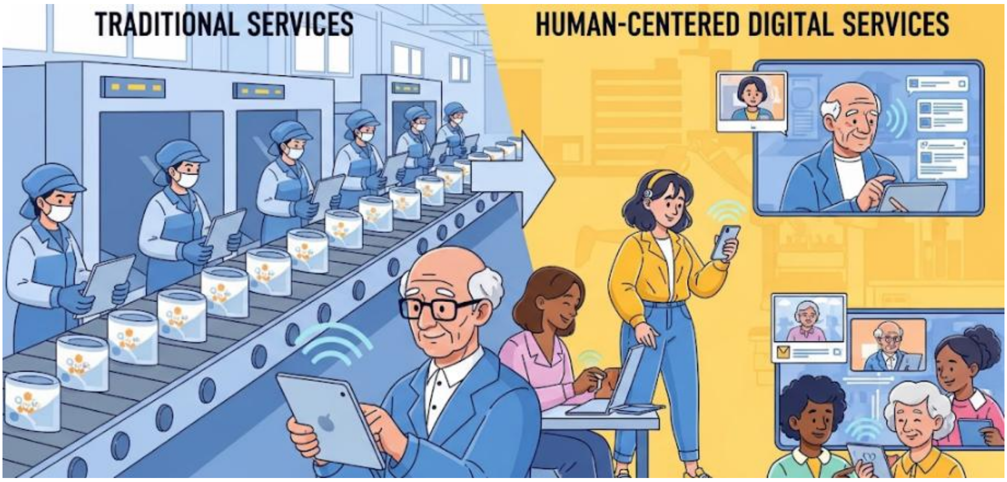

Fig 12 . 1 : Personalized Services in the Digital Age

12.2.1. Definition and Importance

We perform a variety of work for others — in settings ranging from volatile environments of service-oriented interactions to the more stable constructs of product-oriented marketplaces and factories that support the economy. Many of these interactions occur in the context of demanding timelines and a temporally distributed work process that is centered on a specific purpose, such as driving from point to point, executing a construction or manufacturing task, or providing a financial service; such work is a service. More importantly, this work must continually adapt to fluctuating user values of quality, cost, and time; changing situational contexts that require more flexibility in meeting user goals; the well-being and experience of users; and underlying evolving technology that strengthens the service. This definition implies that all services are conjoined with high levels of supporting social interactions directed by varying specific situational constraints.

Such a focus on the changeable needs, goals, and well-being of users — the so-called demand side of services — is a prime consideration for the notion of human-centered services, which thus emphasizes the role of the user. There is also a supply side that emphasizes organizations and managers, the motivations of work for others, service firm performance measures, and innovative growth strategies directed at accomplishing the sustainable organizational objective of satisfying consumers' needs. A service is defined as an act or series of acts performed for another party; this other party is the primary consumer. Their role assumes important input, feedback, and observed output provisions through the evolution of the service process, casting the consumer in an important role for quality control and assessing satisfaction.

12.2.2. Historical Context

Historically, human-centered services emerged much earlier than in the digital area and were also part of electronic servitization. Despite society's current dependency on digital technology and computer capability, enabling a new class of autonomous intelligent systems to play a major role in supporting the next generation of HCS is by no means novel. Of course, historical precedent is not destiny, nor is it a true narrative that standing on the shoulders of giants has informed speculation, or that greater HCS is imminent because HCS occurred before. Nonetheless, the right lessons learned can inform development and deployment towards our future enlightenment. Therefore, we explore the forays into self-service delivery systems and discover what early research suggests guided initial quality evaluation, design, implementation, and even research strategy. In so doing, we discover a history of multiple HCS research deployment cycles, which have occurred in concert with a parallel gradual change to modern demands, expectations, attitudes, and ultimately design configuration for digital HCS.

The first applications of self-service HCSs can be found in the early 1970s. An early ATM model made use of an invented ATM insulator to allow successful HCS cash withdrawals to be accomplished by customers in a matter of seconds. The need for convenience and accessibility, especially outside of banking hours was at least as important as reduced transaction costs. However, that need was manipulated by branches usually not being available outside of traditional banking hours, with banks being overstaffed during the day. It is no coincidence that the first email system allowed users to pre-process mail in the form of composing and then holding until release favorite messages for delivery, rather than needing to be presented with one at a time in real-time.

12.2.3. Current Trends

Elderly care services, education services, and social service agencies serve humanity by helping others and enabling them to overcome adversities in their lives. Such organizations do serve a purpose in-line with long-term benefits for humankind. However, the benefit of such organizations is diffused and not measured in terms of immediate monetary returns. Nevertheless, these services come with their complexities in delivery - identifying needs, coordination of activities, assessment of progress, addressing deterred response, etc. Add to that layer of complexity, internal organizational challenges, changes in budget funding, and changes in the socio-political priorities of the respective government. Also staffing the services is difficult with high employee turnover and inconsistent staff qualifications. These complexities, coupled with a reduction in available resources further compound the human factor service performance, effectiveness, and sustainability.

In spite of the considerable challenges, there is extensive interest in traveling where novice expertise resides to experience services being delivered. Trends show a variety of new applications for mobile client self-management of services being deployed. Cloud technology and expert algorithms are reducing the cost of solutions to enhance outcomes with less adverse risk especially where obvious customer benefits are returning. This has resulted in lowered organization operational costs and extended capability operations. Furthermore, improved services are being delivered with limited budgets by moving on-site and working with citizens to co-design the processes. Important in sustaining this momentum are recognized national champions and continuous re-engineering of service delivery processes. These assistance-enabling services do not solve the identified problem, but provide the temporary means for better person capabilities to cope with adversities in their lives.

12.3. Autonomous Decision Support Systems

We refer to systems that are expected to act autonomously for decision support in complex environments populated by humans and other autonomous entities, benefit from intelligent service capabilities provided by heterogeneous AI components, and be seamlessly integrated into existing socio-technical infrastructures and service ecosystems as A-DS. The role of A-DS is to enhance the collaborative nature of decision-making, offering collective intelligence to groups of decision-makers that are human or autonomous. A-DS sits at the intersection of the human-centered AI and the smart autonomous components technical sectors. Typical examples of A-DS include teams of autonomous agents that perform missions with the assistance of human decision-makers and mission managers, such as in the case of an urban search and rescue, holographic digital avatars or smart immersive assistant environments, whose mission is to facilitate workers in their day-to-day operations by helping them in executing complex workflows. A-DSS also includes smart virtual assistants, dedicated to assisting workers or ordinary citizens in routine operations, such as scheduling meetings, booking hotels or ordering food, and interacting with the corresponding user through speech and text. Other examples include groups of collaborative robots that help human co-workers in manufacturing tasks, from feeding the assembly lines to collaborating in the delicate operations of assembling small devices. Other examples derive from the development of specialized software applications based on a combination of various AI capabilities that are capable of helping crowds solve exceptionally complex decision problems.

12.3.1. Overview of Autonomous Systems

Decision support systems (DSS) have evolved extensively through improved integration with communication and computation technologies. Notably, the last three decades have witnessed a shift from off-line, design-oriented support tools to online and flexible cognitive and affective moderator roles, and from purely internal organizational operations to external and large-scale operations responsive to the organization's external environment. This evolution has sparked interest from researchers, practitioners, and managers given the growing complexity of the environment in which organizations operate, as well as in which autonomous systems are deployed. Indeed, DSS has been transformed from automated process control to providing a dual-user interface with autonomous systems and from managing purely internal operations to also supporting external relationships of creation and co-action. However, there is still little understanding of how DSS will enable the provision of a new generation of human-centered services, with a particular focus on autonomous systems such as autonomous vehicles, robotics, virtual agents, and crowds. In parallel, growing volumes of autonomous systems demand increasingly diverse and intricate components of the decision-support process, for which collaboration between humans, and autonomous systems can affect the success or failure of the overall support process.

Given the fast-paced technological advancements of autonomous systems and the increasing interest in understanding their impact on organizations and society, there is a need to establish a conceptual framework that can integrate these different trends in the research and practice of decision support research. In particular, autonomous systems are a collection of physical and virtual systems that deliver stand-alone services, or members of heterogeneous teams supporting both simple and complex human activities and interactions. Embedded in the digital ecosystems that we live in today, autonomous systems operate increasingly for, with, and alongside humans. Such systems leverage and provide assistance services that enable humans and other teams of autonomous systems to become more productive, creative, and adaptive. We can categorize the functionalities of autonomous systems according to the four stages of the decision-making process. For these systems to be truly autonomous, they must be capable of performing these functions without any human intervention or supervision.

12.3.2. Technological Advancements

The fifth wave of applied artificial intelligence, where AI-driven systems are entrusted with higher levels of autonomy in complex realizations and user populations, is enabled by the convergence of five technological trajectories. The first trajectory is the advance of the available sensor and actuator technology covering a broad spectrum of modalities including vision, audio, touch, and motion for inputs while haptic interfaces, AR, and VR systems for outputs. The second trajectory is the fast-growing computational and communication infrastructure which consists of high-performance and efficient chips and distributed heterogeneous edge-cloud systems with 5G. This infrastructure both enables deep learning computations for sensor information abstractions and supports the connectedness of autonomous decision support systems and the communities of users and collaborators. The third trajectory is the advance of the representation models for learning to calculate diverse decision-relevant abstractions from the raw sensor signals, including deep learning architectures of higher performance and generalizability, as well as capable semi-supervised and unsupervised training paradigms. The fourth trajectory is automated learning for customizing decision-making systems to specific users, tasks, and contexts from limited supervision utilizing both classical structured prediction methods and approximate gradient and inverse parameter learning techniques. The fifth trajectory is the design, validation, and trust of human-centric autonomous decision support systems, where explicit consideration is given to the impact on the users and user populations of the data-driven learned solution.

12.3.3. Applications in Various Domains

Such autonomous systems - so-called artificial companions, and interfaces can be of various types: like avatars or socially aware, robot-like systems; they can be specialized in distinctive modalities or utilize diverse modes of interacting: like industrial or gaming robots, or conversational, multimodal, or cognitive assistants. The domain of cognitive or conversational assistants, like similar mobile phone functions, or conversational agents in the form of dialogues in chat interfaces, is naturally the first arena of deploying self-learning autonomous assistants. Their capability to develop various competencies on the client's data, and contribute to accelerating various activities due to their particular knowledge of the client's needs and styles, is evident. For now, they are generally hybrid, making use of self-learning technology for specialist tasks but relying on template-based approaches for the basic competencies. Such services are frequently created within creative partnerships with clients or particular communities, utilizing methods that conform to the new semiotic paradigms. With the cognitive assistants being developed gradually also end-users become more knowledgeable about deploying them for their specific tasks and needs, which creates an environment suitable for facilitating the automating tasks and co-creating activities with the machines. Autonomous decision systems can also be utilized for producing solutions for numerous personal or collective needs, instead of merely supporting and advising users in their choices. Autonomous partners manage conflictual or cooperative collective plans where the achievements of a person, group, or society may or may not be synchronized. Such abilities are evident in the case of games, a tradition that started a long time ago, but today's expert partners are mostly specialized. In professional or personal settings autonomous partners in artists' or athletes' roles can be utilized for realistic simulation and training, or they can act as avatars for creating peculiar semiotic products.

12.4. The Intersection of Human-Centered Design and Technology

We languagize and encapsulate the world to orient both users and deployers of technology, yet fail to enable meaning. We develop and introduce innovations that have little impact on our inhabitants, despite them being seen as "engaged" or "connected". This trend of proliferation of lifelike interactive services wrenches life and meaning from the human equation. Using the devices, their pseudogenes, interconnected personalities, disembodied representations, and electronic auras homogenizes the superficial at the cost of the human. Radical advances in sensing-based, context-sensitive, and intelligent technology have contributed to an immense diversity of applications that support our activities, but the challenges at the intersection of human-centered design and technology remain.

There are established principles and guidelines for human-centered design, developed over nearly half a decade of research and development. These include understanding and interpreting how people experience the world, interpreting their needs, translating those needs into meaningful and emotionally satisfying products and systems, developing user-friendly accessibility solutions, and performing satisfaction testing before launch. Yet, the guiding focus of these principles and guidelines is on the products themselves, not on the larger, collective systems of heterogeneous "things" on which we depend for basic governance and services like security, health care, education, energy, and transportation. Also, today people experience their world through a seemingly infinite physical and electronic infrastructure that is entirely disjoint, unmanageable, and largely unavailable at the right time and place. It is asynchronous, insensible to context, and usually hidden from our view. Each of our interactions leaves only a momentary mark on the electronic version of ourselves, if at all. Activities of interest are little more than prosaic, erratic movements in the stochastic background of collective electronic consciousness.

12.4.1. Principles of Human-Centered Design

A principal tenet of Human-Centered Design (HCD) is that the final design must fit the needs of the people using it. This paper offers answers to this question; not a complete answer, but new steps towards a more complete answer. HCD has gained attention in both research and practice in recent years. In particular, many new questions have arisen about the details of HCD. The principle of putting humans at the center of design needs to be more clearly enunciated and reflected in systems being created by designers. HCD is summarily said to be both an approach to practice and the outcome of that practice. Putting people at the center of the design is a deceptively simple statement; the challenge is to elaborate it into what it means to HCD as an approach to practice and the outcome of that practice. Here we identify three key conceptions of the principle of putting humans at the center of design; these are usability, stakeholder involvement, and user desirability.

The first conception is usability, which has been formalized in the HCI community as accessibility, utility, and learnability. Accessibility requires that people with a range of abilities, and disabilities be able to use and benefit from designed systems. Second is the concept of stakeholder involvement in design, which seeks to include a range of users in the design process. HCD should consider all those affected by the design, not only a subset chosen for business purposes. These different stakeholders contribute both different perspectives and different requirements to the design. Third, is user desirability, which argues that design should consider the desires, goals motivations, and emotional responses of users. User desirability adds a more emotional dimension to design systems, and provides user delight, attractive products, and evokes a more positive motivation toward interaction.

12.4.2. Integrating Technology with Human Needs

The design of healthcare interventions is often only driven by constraints applied by funding bodies, service delivery contracts, strategic systems, or clinical needs. In many ways, the pragmatic users of those services — patients and their families — are often a secondary factor in the design of healthcare systems. Policies that govern eligibility for healthcare services and frameworks for service delivery dictate the way services and systems are designed. How burdensome the processes involved may become, or even how off-putting the services are to those patients in need, are rarely considered until long into the development of a healthcare service or system. This can lead to costly delays in the trials, evaluations, and roll-out of new services. Technology can transform how healthcare services are delivered. Neuroimaging has changed the understanding of how the brain works and treats those with conditions such as depression and anxiety. Minimally invasive surgical techniques have revolutionized surgery for a number of medical conditions. These examples show how new technologies can break through previous limits of how things are done, and how much we understand about a certain condition.

The key distinction with the HCD perspective is the incorporation of the patient perspective into the development of those tools, systems, or standards being set. Without that patient perspective, there’s a risk of developing a technology — or other systematic tool — that, despite being sophisticated, effective, and well-advanced, fails to consider how that technology will impact those affected. The priority in any intervention, however it might be classified, is the needs of the patients undergoing that intervention, and technology must be used to enable service delivery that aligns with patient expectations. If not, there’s a genuine risk of moving healthcare delivery into the realms of an artificial construct that may have little relation to the real-life experiences of the patients using those services. No amount of technology sophistication can overcome the fact that these are real people undergoing suffering.

12.5. Ethical Considerations in Autonomous Decision Making

Within the ever-changing landscape of the development and integration of autonomous decision-making applications into human-centered services, the importance of tacitly embedding ethical considerations into such applications is becoming increasingly visible. Concerns around fairness, transparency, accountability, privacy, and the need for appropriate human oversight become essential to explain and model application development. Crucially, it is necessary to recognize that the above alone does not ensure that such systems are indeed beneficial to humans, just as increasing levels of automation in decision-making processes are not guaranteed to have positive outcomes. It is here that meaningful and appropriate interaction with users and stakeholders during the design process, and afterward, becomes informative of the veracity of such technology across a range of criteria designated to assess levels of user trust and reliance in proposed human-AI collaborations.

Bias and Fairness

It is broadly recognized in the Social and Data Science communities that the data-driven learning of prediction and recommendation algorithms is at risk of amplifying existing biases in the input data, especially if they stem from historical decisions made by biased actors. These concerns are especially highlighted in the context of deploying AIDAS in sensitive domains such as justice, finance, and hiring, where feedback loops resulting from increased reliance on biased algorithms in societal decision-making have the potential to cause harm to the demographic groups already at risk of being marginalized by biased stereotypical decisions and their associated outcomes. Bias mitigation, a.k.a. fairness-aware Learning, has been thus established as a priority area of research within the AIDAS community. The related training fairness objectives and additional fairness constraints are aimed at minimizing bias amplification and, in relation to the decision-making applications for which the AIDAS are deployed, stipulate desirable fairness conditions such as demographic parity or equal odds, often given different groups defined in terms of sensitive attributes such as racial or ethnic origin.

12.5.1. Bias and Fairness

The idea that Artificial Intelligence (AI) algorithms are objective and unbiased has long been challenged. The algorithms we create are trained on biased training data that reflects and perpetuates the unfairness of society. Unfairness has practical consequences for organizations that deploy AI algorithms, such as unreliable referrals in hiring and negative public perception. Given the intent of Autonomous Decision Support to create user-centered services at a small scale, identifying and expelling bias from decision-making processes is a prerequisite. In particular, the underlying burden of biased algorithmic decision-making in sensitive areas like hiring, finance, or credit retrieval must be recognized from the outset. Developers involve stakeholders early on in the process such that clear success disclosures are identified, user expectations are aligned, and possible adverse effects are mitigated from the outset. The technical means to remove unfairness hinges on setting feasible fairness standards for the specific problem. Fairness definitions differ: Do we want all demographic groups to receive the same rate? Or do we want all demographic groups to receive the action that is predicted to be best for them? The two definitions lead to different results: group fairness dismisses that individuals vary, and individual fairness does not ensure fairness on the group level. The preferred definition dictates which technical means need to be employed to limit group or individual respectively.

12.5.2. Transparency and Accountability

Introduction

Organizations employing the use of autonomous decision services should proactively prepare for inquiries in the event they make a decision that the customer is not satisfied with. These inquiries can take the form of required audits and assessments by government agencies, or inquiries from the customers themselves as consumers become aware that a decision made on their behalf was not made using a human decision service. One means of answering these inquiries, thus achieving accountability and transparency for actions taken on behalf of the consumer, is to emphasize what checks and balances prior to a decision and through the authoritative use of a decision service, were made that ensure the decision is indeed in the best interest of the consumer. These measures also serve to maintain the trust of the consumer in such services and in turn the popularity of such services as they are offered.

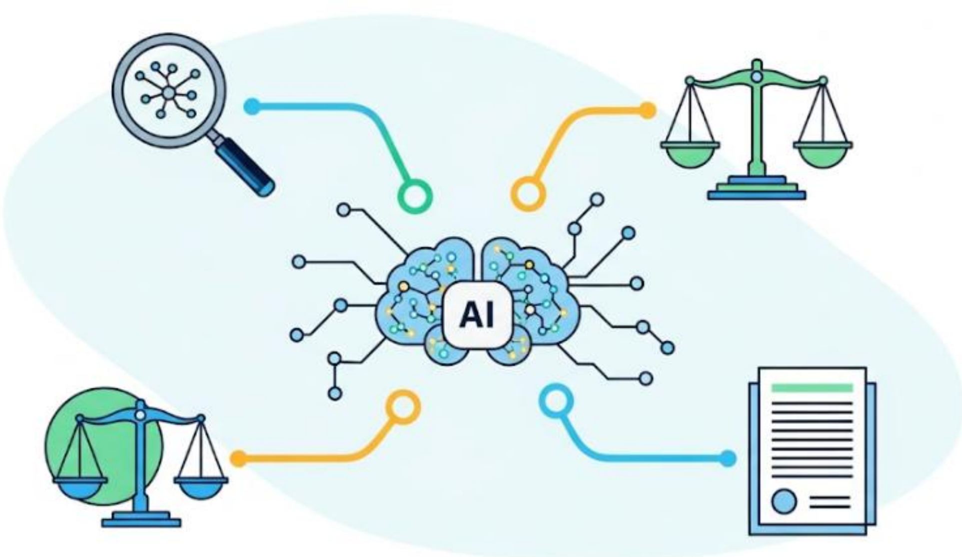

Fig 12 . 2 : Transparency in AI Decisions

Internal Organizational Remedies to Decision Service Transparency and Accountability

There are measures that an organization can take that will preserve consumer trust in a decision by a decision service, and thus provide accountability if a customer is dissatisfied with the decision made by either an autonomous or human decision service. The purpose of such measures is to provide transparency and a means of inquiry in the event a consumer is displeased with a decision made on their behalf and to enable accountability within the organization for actions taken on behalf of the client. One measure is through auditing and testing, as discussed. The organization can demonstrate through selected data that a decision model and service is performing appropriately, or not. They can also explore the behavior of the model on or off their data through testing. This can be further enhanced by the origination and use of documents that lay out how the decision model was created, and tested, and how it is maintained. These documents can serve as a guide for consumers as well as regulators in the use of a particular decision service through a particular model, effectively transcending the model itself.

12.5.3. User Trust and Acceptance

The aforementioned factors of bias, transparency, and accountability lay the foundation for user trust and acceptance which are important for the future use of Autonomous Decision Support. Trust is a belief in the efficacy and benevolence of an actor and guides our willingness to rely upon their offerings during decision-making. While traditional decision-support systems help people in their decision-making, Autonomous decision-support systems such as machine learning, recommender systems, expert systems, and even predictive algorithms are taking over so much responsibility that humans become mere consumers or resources in the decision-making process. Especially in areas where rejecting advice might lead to damaging consequences, people rely heavily on Autonomous Decision Support Systems triggering questions about user acceptance and trustworthiness. As users increasingly rely on algorithms to guide decision-making, algorithmic explanations for decision outcomes may impact the user's perception of an algorithm's trustworthiness and explainability. Without such explanations, many users may view algorithmic recommendations with caution and skepticism, leading to diminished usability of algorithmically based recommendation and staffing systems. In order to enable efficient cooperation between humans and machines, it is important to foster trust in the systems.

Trust in an Advisor is determined not only by how accurate its predictions are but also by the advisor's past inconsistencies; the reputation of the advisor; its explanations of difficult predictions; and its recommendation strategies. Adopting the use of recommender systems for more sensitive areas of decision-making such as criminal justice, health care, or hiring will only succeed if the users of such systems can identify with them. It is important to establish these systems with the users in mind and focus on aspects of inclusion and diversity. However, trust is intrinsically linked with perceived usefulness and perceived ease of use.

12.6. Challenges and Barriers to Adoption

The success of the next generation of autonomous services lies not only in their technical prowess but also in their widespread acceptance by both users and providers in society.

Such services can only be enhanced if uptake is high, and the friction in people's lives that they seek to remove is diminished significantly. The advantages of integrating such solutions into society can also be diminished if a longer time is taken to build confidence in the capabilities, tools, and systems adopted. As solutions become more commonplace, concerns begin to recede and initial anticipatory anxiety dampens. Such delays may also mean systems and solutions become embedded offering a less rich opportunity to capitalize on the novel attributes the solutions offer.

Despite the potential, obstacles remain to the successful integration of autonomous solutions into people's lives. Some are prohibitively expensive or impossible to reduce to practical reality; others are simply a matter of developing the awareness of services and the trust in them, be that from a consumer or provider perspective. Other barriers include the speed at which such solutions emerge, the evolving, complex rules about regulation, responsibility, and liability; concerns about job replacement in some sectors, and opportunity availability in others; privacy, security, and misuse apprehensions; and a slow systemic change approach stifling innovation and skewing the potential benefits, balancing exploitation and exploration, for economic gain. These topics, and a summary of the component areas, are now discussed.

Although developing technical capabilities is a concern, further highlighted by the current pandemic, the uptake of solutions also needs addressing to ensure the sustainability of the new service ecosystem. Lack of capability in certain areas of society is prohibitive, especially across older populations, and need not be cause for concern. The focus should be on developing ethical services using transparent systems that naturally facilitate the uptake of a capable service.

12.6.1. Technical Limitations

When autonomous decision support capabilities approach the level necessary for societal relevance, there will be two general concerns about their accuracy. The first is whether, within the domain in which it operates, the performance of the system is so poor that it is worse than current human performance. Second, if such systems exceed current human performance, for what classes of decisions is that the case? The first point is often the focus of concern; investment into human-equivalent performance seems ill-advised. However, it could be that only a subset of common tasks is completely mismanaged and that there are sufficient other tasks that are accelerated by humans (and supported by autonomous decision support tools) that the industry would adopt such a system. A reasonable approach is that autonomous decision support need not exceed current human performance, but should take a significant fraction of decisions out of the human feedback loop either way. Measuring that performance is a research challenge. When the accuracy of autonomous decision support is significantly higher than that of a human, we could use set-wise accuracy to compare them on some scalable set of problems. Second, if the human is better, we can use that function when developing or tuning the decision support function.

The second issue is that the level of autonomy may vary, and again there exists a need for calibration. It would certainly be possible to have calibration systems that dynamically tune the level of autonomy upwards or downwards based on feedback from the human operator. But for such a solution to enable acceptable human-machine teams its calibration must be trustworthy, and that may be an open research question, particularly because its behavior will vary based on circumstances, such as the amount of time remaining for the human to be able to complete the task. Furthermore, there are further problems in ensuring the trustworthiness of the system as members of the team differ. Addressing these challenges is an area of ongoing work, but work to develop accurate (or highly calibrated) autonomous decision systems will aid applications of decision support systems that enhance human capabilities and allow teams of humans and machines to work productively together.

12.6.2. Cultural Resistance

The implementation of advanced autonomous decision support services such as AI content generation or public analytics presents hurdles in the context of governance because governments are blamed for decisions and creating a resistant mindset towards the adoption of automated assistance. More efforts for overcoming obstacles and hesitations towards integrating the technology to facilitate operations and service management would be of use. Despite the potential anticipated benefits of technology adoption, the degree of reluctance to change among various actors is overwhelming. Prejudices and potential bottlenecks would need to be overcome before performant agencies fully embrace the use of available tools that improve operations, enhance accountability, and strengthen organizational development. Policymakers have raised concerns about the accuracy and disinformation risks with the advent of AI tools and resist integrating those tools within the agency despite these organizations being architected around the power to act, think, and correctly decode given issues. Yet, the development and ethical considerations still could not dismiss the inherent risk carried by the responsibility to create, modify, and evaluate the content. For some, automation could entail a disconnection from the social world or the contemplation of “creatures” designed to experiment with language and social interactions. Theoretically, any device that produces a grammatical output associated with reliability and accuracy at some level is right at the boundary of our common definition of obligations; perhaps obligations aren’t the intrinsically defining feature of human activity. The relational facet also stands out within our plethora of criteria for framing the technology usage around the public sphere and establishing cohesive and harmonic policy communication dependence.

12.6.3. Regulatory Issues

Non-specialist users wishing to augment their decision-making capability with the insights and experiences captured inside decision support systems, in the form of models, tools, and visualizations, need a pragmatic way to understand the implications of the decisions they are making when deploying these systems at work. However, the feedback individuals receive on the nature and scope of these implications by virtue of their working in regulated environments can impact the adoption and use of decision-support applications. The more sensitive the data the algorithm is using, the trickier these challenges become.

There may be no explicit regulation about the internal use of a decision support application that has been developed internally by an organization. However, individual employees may have certain rights and expectations around privacy and confidentiality, especially when using a tool that makes them expose their personal information when performing certain HR processes or offering support for a mental health issue. Their decision to expose themselves and seek help through a decision support application should have a similar level of protection as if they were seeking help directly from physicians or psychologists. System designs should also consider the responsibility of meeting the legal implications of user-driven decisions when using regulated decision support systems. These implications will undoubtedly vary depending on the type of system being researched and the specific legislation that applies.

12.7. Future Directions for Research and Development

Research and development (R&D) are key to innovation and the growth, efficiency, and productivity of knowledge-driven economies. New products and services increase standards of living and directly contribute to the wealth of nations. While this is a reflection of work done over the previous years, it is clear that further advances in capability delivery for the development of work are required. Autonomous decision support is culminating from substantial innovation in services across a plethora of domains. Thus, how do we advance the R&D investment in autonomous decision support to enable the broader range of applications that are potentially realizable? Such as those that bolster government and create the environment to stimulate the economy and grow sectors such as security and health. Ultimately, such R&D investment has the potential to advance services to encapsulate the vision of a democratically accountable and technologically advanced society.

Some further questions arise. What emerging technologies exist that can contribute to the fulfillment of the pathways to citizen-centric services? What existing, and what additional, collaborative models and joint, interdisciplinary approaches are required to further the exploitation of these emerging technologies? We would suggest that some, new, focus is needed on both the technologies and the joint, collaborative aspects of doing R&D that result in the fast realization of its promises. We would suggest that the greatest advances in new product and service development accrue from collaborations both in the business and tertiary education sectors. A particular focus of collaborative efforts needs to be on the transfer of knowledge and insights gained in new technology and capability developments into implementable systems.

12.7.1. Emerging Technologies

Technologies such as Blockchain, Naturally-Aware User Interfaces, Wide-bandwidth Haptic technology, Wireless Body Area Networks, and Ambient Intelligence are already exciting research and commercial movements. These and other technologies will continue to advance, become cheaper, and saturate the environment. They will help introduce and deploy many new sensor nodes that collaborate with internal world models and become coupled to the human user and with each other, Grounded Intelligent Agents. Such nodes will be supported by the continuously growing sensor web, collecting and indexing information. Paradigmatic research and development should result in the next generation of human-centered and moored computational services being created by others based on plug-and-play systems of Grounded Intelligent Agents cooperating locally at an increasing number of places and in an increasing number of expanding domains. The ongoing trend of IT services becoming completely modular and choreographed by non-technical, end-users enhances the potential of such services and their variety. It also means, however, that Grounded Intelligent Agents should introspect what support instruments to make easily available through what interfaces for non-technical users.

Grounded Intelligent Agents should also implement and communicate affordances for other agents and be integrable and choreographable at all levels of the hierarchy by other agents so that seamless interaction among, coordination across, and choreographed, cooperative coupling with and also between Human Cognitive Systems becomes supported automatically, becoming Intelligent Environments. We envisage Intelligent Environments becoming specialized in automating the end-user, be it a technical expert or not, into the cognitive services integration and choreography phase toward Intelligent Environments. This assumes that Intelligent Environments have no or very little symmetry breaking available.

12.7.2. Interdisciplinary Approaches

In this section, we present what we see as some guidelines for those researchers interested in the interdisciplinary development of models of autonomous decision support, either as a preparation for automated support for human-centric services or as areas of application for the implementation of such systems. We also highlight possible paths towards the establishment of scientific disciplines devoted to the study of the type of decision models that would make possible the deployment of reliable systems for services that are complementary to human action and decision-making.

In our framework, the study of models of decision support does not only need to consider possible improvement in parameter tuning processes optimized by generic machine learning methods but should focus on the relation with specific human–decision models, to whatever degree or through whatever method those models can be realized that come from the disciplines that study them, such as UIS, learning, social behavior, etc. Also on the topics considered by specific social and natural sciences. This will imply the possibility of transfer methods and findings from those disciplines where such expertise exists. Collaborations between engineers, normalization, modeling, and domain experts should be conducted as soon as possible in the specific interdisciplinary teams involved with specific research aims. Collaborative projects between universities accounting for the different fields involved and specific industries working in fields of common interest could constitute a significant impact on the practical experimentation of joint methodologies coming from distinct scientific fields.

To summarize, we have considered models for autonomous decision support destined to be used in human-centered work and to allow for decision automation by AI of functions presently limited to humans in the context of services to people as key enablers of the next generation of human-AI collaboration.

12.8. Conclusion

Our vision is to pioneer a new generation of intelligent services that optimize the ability of people to support each other while helping those who need help most, fostering transparency, empowerment, and well-being. We believe that, through deploying autonomous decision support, peer endorsement, and open data, it is possible to put the next generation of human-centered services back towards their original intent - build resilient, self-sustaining communities that can help each other, while also being backed by decision support technology that can augment the often difficult, emotionally-charged, and complex nature surrounding decision-making.

We have identified the need for highly enabling decision-support services, grounded on capabilities that help humans to do what humans excel at - foster community and social bonds - while using technology to support and optimize this process. Together with open data, data sharing, and endorsement, we believe that this will ensure a new generation of services that are accountable and empowering.

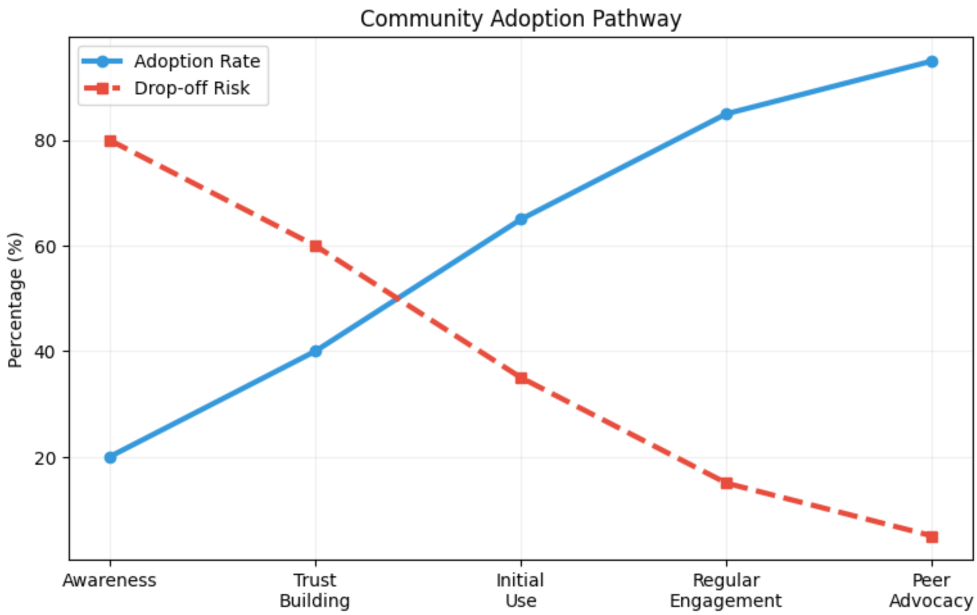

Fig 12 . 3 : Community Adoption Pathway

12.8.1. Summary of Key Insights and Implications

This essay presents a vision for the next generation of human-centered services that transcend elaborative, reactive decision support for complex services, such as customer-care services, to enable more effective autonomous decisioning for a wider range of human-centered services, such as e-government services. This is enabled by a framework for cross-service interdependencies that has several significant positive boundary-control implications and is made possible by advances in natural language processing, semantic knowledge representation, cloud computing, data analytics, and integration technologies. To place this vision in context and provoke indeed the discussions in the future sessions, we conclude with a summary of the key insights and implications underlying the vision.

Services have traditionally been characterized by their unique capabilities: for economic and social value; for facilitating exchange; and for linking participants. Decision support for service encounters has also usually been elaborative, reactive, and batch-oriented.

The enabling of choreographed and then autologous coordination is a major advancement for decision support for services. Autologous decision coordination, then of actually delegating key options to the service consumer requires more complex capabilities but uses the same advancement. Service systems are not monolithic, cohort-attribute decision systems, with few members in a cohort of decision-need volatility. Other decision systems in the same service environment might, and often will, enable more accurate needs assessment. Autologous decision coordination enables directive service, e.g. recommending camera X for consumer Y, based on her profile and purchase record, as well as elicitive service, e.g. asking consumers for their preferences.

An essential, but non-instrumental, part of both is the communication of the decision system member's intentions to other service system members. Autologous decision coordination can also enable more seamless customer journeys that span not only but also outside the service-enabling touchpoints. These and other positive implications can only be realized if the service consumer is willing and able to delegate some service-related decision-making to the service provider.

References

Endsley M.R. (2017). Autonomy and Human–Automation Interaction. Human Factors, 59(1), 5–27. https://doi.org/10.1177/0018720816681350

Kamar E. (2016). Directions in Human-AI Interaction. Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, 6, 4291–4296. https://doi.org/10.24963/ijcai.2016/603

Susskind R., Susskind D. (2015). The Future of the Professions: How Technology Will Transform the Work of Human Experts. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198713395.001.0001

Gajos K.Z., Czerwinski M., Tan D.S. (2018). Human-Centered AI: Principles and Practices. AI Magazine, 39(4), 19–28. https://doi.org/10.1609/aimag.v39i4.2822

Smith, P. J., Geddes, N. D., & Beatty, R. (2009). Human-centered design of decision-support systems. In Human-Computer Interaction (pp. 263-292). CRC Press.

Nguyen, H. N., Lasa, G., Iriarte, I., Atxa, A., Unamuno, G., & Galfarsoro, G. (2022). Human-centered design for advanced services: A multidimensional design methodology. Advanced Engineering Informatics, 53, 101720.