

Chapter 2

SmartData: The Need, the Goal and the Challenge

George Tomko

Abstract Technology must form an integral component in the defense of our personal privacy. Policies and regulations will serve, at best, as lagging remedies in the fast-paced world of cyberspace. In a world where personal information can increasingly be transmitted and used in multiple locations simultaneously, protecting privacy may only truly be accomplished if the information itself becomes “intelligent” and capable of making appropriate decisions relating to its release, on behalf of the data subject. In other words, the data must become smart—we need SmartData. This presentation will discuss the growing need, the challenges, and ultimately, the benefits of developing intelligent agents to protect our privacy online.

2.1 A Recipe for Social Decline: Expand Surveillance and Shrink Privacy

New scientific discoveries and innovations in technology are the lifeblood of society’s well-being and prosperity. This lifeblood depends on a continued political and cultural context of freedom, which like oxygen, supplies the energy for innovation. It is not coincidental that the western world, which enjoys the most political freedoms, has also become the most innovative and prosperous. However, the pillars of freedom, which include respect for privacy and protection of individual and property rights, are being jeopardized in the pursuit of public safety against criminals and potential terrorist threats. In this pursuit, society is experiencing a greater expansion of electronic surveillance, increased misuse of personal information, and the concomitant erosion of civil liberties. In its response to terrorism, mankind has excelled at developing technologies of surveillance against an

G. Tomko (✉)

Identity, Privacy and Security Institute (IPSI), Expert-in-Resident,
University of Toronto, Toronto, ON, Canada
e-mail: gjtomko@hotmail.com

ever-expanding list of security threats identified by governments around the world. The concern, though, is that as these technologies become more sophisticated and incorporate recent advances in artificial intelligence, they will become a threat not only to our nations' enemies, but also to its citizens. Loss of freedoms is always precipitated initially by loss of privacy before remaining individual and property rights are ravaged.

This trend of surveillance is repeated in the private sector, where corporations are also collecting more personal data about their customers in their quest for additional revenues and marketing advantage. This fact alone poses privacy challenges, but now governments are "encouraging" corporations to share the customers' personal data that they have obtained through business-related transactions in order to build personal profiles and identify potential terrorists or criminals. Governments are promulgating a widespread view that in order to protect citizens against the new threats of the 21st century, freedoms such as civil liberties and the right to privacy must be relinquished to some extent. They espouse a zero-sum paradigm wherein public safety may only be protected at the expense of our freedoms, especially privacy.

Not only is this view especially flawed, but it is dangerous. It is flawed in that it demonstrates a fundamental ignorance of technology and the power of innovation. There are far better ways to design information systems which afford a positive-sum outcome rather than adopting a zero-sum paradigm which inevitably leads to ongoing reductions in privacy when pitted against the need for security. This is the basis of *Privacy by Design*—building into technology the capability to achieve multiple functionalities—public safety and privacy, or using personal data within the constraints of privacy such that both businesses and users may benefit. Adopting a zero-sum paradigm is dangerous because curtailment of our privacy and freedoms will ultimately stifle innovation, lead to mistrust and fear of our governments and corporations, and diminish the prosperity of our society. It will also dampen the joy of participating in society—reflect upon how much fun it is to travel by air these days!

2.2 And Now for Something Completely Different: SmartData!

We believe there is a better way. In the spirit of *Privacy by Design*, that way is to use artificial intelligence to protect privacy, civil liberties and public safety—to build "SmartData"—data that protects itself in a manner that is sensitive to the needs of the data subject (to whom the data relate), while enabling multiple functionalities such as judicially authorized requests for data.

SmartData is a novel technique that empowers personal data by wrapping it in a "cloak of intelligence." It is an intelligent agent that becomes the individual's virtual proxy in cyberspace, controlling the release of his or her data, in accordance

with the user's preferences and the context of the situation. As such, no longer will personal data be stored in the cloud or shared as merely "data," encrypted or otherwise; it will now be stored and shared as a constituent of the binary string specifying the cognitive structure of the SmartData agent. This binary string when downloaded into a processor resembling a reconfigurable computer will configure the processor to embody a cognitive structure which will proactively protect privacy and security, right from the outset, so that it is not treated as an afterthought. It embodies a foundation of control and trust within the technology itself as the first line of defense, incorporating the principles of purpose specification, personal consent, security, and use limitation—the essentials of user control. The SmartData project incorporates the advances made in the technology of simulating virtual worlds, together with the ideas emerging from the fields of machine learning, evolutionary robotics and embodied cognition, within a framework of dynamical systems as a methodology toward this ultimate goal.

2.3 How We Got Here: The Time for SmartData Is Now

Prior to digital databases, individuals had far greater control of their personal data, since by and large, their consent was required each time it was released. Furthermore, any personal data that was released had an effective half-life—a function of human memory. With digital databases however, this control has been lost since the individual cannot be present at all locations in cyberspace where his or her personal data may likely be used. Furthermore, personal data no longer has a half-life, now remaining largely permanent and available indefinitely. It is now the organizations governing the databases that have effective control over its uses. Although subject to privacy laws, the complexity of technological infrastructures and cross-border regulations makes it increasingly more difficult to "police" adherence to such laws, and once a data breach has occurred, the personal data and its potential unauthorized uses can no longer be controlled. Therefore, "regulatory compliance alone is unsustainable as the sole model for ensuring the future of privacy" [2].

Privacy by Design in its initial form (*PbD* 1.0) was an attempt to have organizations develop a methodology, which included policies and technological design, to incorporate and build privacy into the systems that handled personal data. But even here, the individual is still dependent on the integrity of the organization to follow Fair Information Practices [1], and the effectiveness of the technology to withstand internal breaches and external hacking. Of greater concern, the trend of governments, under the guise of public safety, to intensify their surveillance of individuals and collect more personal data, has neutered the effectiveness of *PbD* 1.0. Governments can now mandate businesses to give them access to an entire database of personal data under the pretense of public safety without, in many cases, informing the data subjects. This pretense of good intentions is eroding one of the pillars of freedom in our society. *Without privacy, there eventually will be no freedom, and without freedom, there will be no innovation and prosperity.*

SmartData is an evolution in *PbD* (*PbD* 2.0) which moves control from the organization and places it into the hands of the data subject. Instead of having a one-to-many oversight mechanism where one organization is responsible for protecting the privacy of many individuals; the goal is to evolve to a one-to-one oversight mechanism whereby one individual is responsible for one set of personal data—his own. Organizations would no longer have the unfettered use of databases of collective personal data without the express consent of the individual, and furthermore, the individual would be able to, in effect, control the uses of his/her data at all times even when it is being used by organizations for a consented purpose. Ironically, SmartData technology becomes an integral component in the defence of one's privacy that is threatened by misuse of technology itself. In a world where personal information can increasingly be transmitted and used in multiple locations simultaneously, protecting privacy may only truly be accomplished if the information itself becomes “intelligent” and capable of making appropriate decisions, relating to its release, on behalf of the data subject. In other words, the data must become smart.

Ideally, development of a new technology application should solve not only the problem at hand, but also attempt to incorporate solutions that will be effective as technology infrastructure evolves in the future. Recently, Philip Rosedale, inventor of Second Life, made a compelling case that the current flat two-dimensional World Wide Web will be transformed into a 3-D virtual world wherein text and images will only form a subset of the total cyber-environment. He argues that since human beings evolved in a 3-D world and are by nature social animals, a corresponding virtual world would allow more familiar, efficient and social ways of exchanging information. Up to now, however, users have always been “external” to the Web—on the “outside,” looking in—interfacing with the Web directly through a keyboard, or via a computer programmed to carry out instructions. A similar situation exists in current 3-D virtual worlds, such as Second Life, where avatars are, for the most part, directed by the user or a computer-surrogate on the outside, in the “real world.” But this situation is changing with the introduction of agents such as viruses, worms, cookies, and Trojan Horses. Getting “inside” the Web has already started although mainly for malicious purposes. However, these agents are not autonomous. They are essentially “dumb” in that they can only carry out actions based on previously programmed code. Although they have no agency, per se, what is important is that the direction of “agents” moving “inside” the virtual world has begun.

The next evolution in the Internet will be the introduction of intelligent, embodied agents within 3-D virtual worlds, connected to the digital cloud, and having access to a global network of data. These agents, we predict, will become our acting surrogates, thus generating more productive way of exchanging and processing information. Such a 3-D virtual Internet has the potential to inspire totally new innovations, as did the flat Web. However, the need for privacy and security in such a cloud-based virtual environment will be enormous. Although a full blown 3-D virtual internet is still in the future, the introduction of cloud

computing has already exacerbated the difficulty of securing privacy and protecting personal data in cyber-environments, especially when there are competing forces interested in accessing personal information. As mentioned above, governments, public officials and businesses seek unfettered access to such data, for a variety of purposes. On the other hand, consumers generally wish to divulge their personal information only for specific purposes, after which they want the data destroyed. But the difficulty with current data protection schemes, caught in the tug-of-war of competing interests, is that once the data is in plain digital text, it can be easily copied and disclosed, against the expressed wishes of the data subject. Personal information, once released for a singular purpose, may now become lost forever in the “cloud-based virtual worlds,” potentially subject to unending secondary uses.

These difficulties, although tempered by regulatory policies and legislation, can never be completely surmounted because their source arises from the way in which data has existed since the advent of digital databases—as passive in nature, merely bits and bytes on a storage device. At its core, this is the precise problem we are facing: the personal information of an individual, as represented by a binary string of data residing in the cloud, is not capable of protecting itself against unauthorized, secondary uses. In an attempt to overcome potential personal data infringements, a tangled web of international legal provisions and commercial contracts has been established to address the various privacy and proprietary concerns of different jurisdictions. However, the prospect of what we face is not reassuring—not only a global, legal bureaucratic nightmare but also a technical morass of different systems and standards, all trying to interface with each other. Moreover, all of this is overshadowed by the nightmarish prospect of heavy-handed governments motivated by Orwellian “good” intentions, further infringing on our privacy and ultimately our personal freedoms.

While no system can solve all of these issues, we propose that by transforming personal data from a passive string of bits into an “active” form capable of protecting itself, many of the legal and technological issues will be circumvented. A potential side benefit is that the regulatory framework and legal structures between parties need no longer be the first line of defense: they will be transformed into serving as the backstop, in the same way that commercial establishments treat criminal and tort laws against theft of merchandise as a secondary line of defense—with the primary line of defense being “technological”: a secure building, the installation of anti-theft systems, the presence of security staff, guard dogs, and so forth. Unless we are innovative, and are able to solve the privacy vs. public safety issues that arise in a digitally connected world in an analogous technological manner, the innovations themselves may be curtailed since users will not trust the systems, and businesses may refrain from using them. Or far worse, society will creep toward an authoritarian hell along a road that is paved with the seductive good intentions of greater public safety.

2.4 SmartData from 50,000 Feet

The purpose of *Privacy by Design* is to proactively instantiate “Fair Information Practices” into the core of all data-related functions or services on the Web that utilize personal information. To accomplish this, we want to first build a computational foundation for intelligent agents to learn how to protect both the privacy and security of personal information as they go about performing Web-based services on behalf of the user. This computational foundation can then be expanded into other data-related domains such as search, medical diagnostics, electronic health records, social networks and so on, such that all of these new data-related fields have privacy incorporated into their core processing—the essence of *Privacy by Design*.

The three components necessary for achieving SmartData are: (1) securing personal data; (2) embedding data access rules (based on, for example, data subjects preferences, Fair Information Practices, local regulations and potential judicial warrants) “within the agent;” and (3) responding to requests for information contingent on its access rules, background/context and ongoing experience. This is the long-term vision of SmartData. However, our first task is focused on a proof-of-concept demonstrating that the principles inherent in SmartData can be built into a simple Web-app that, for example, orders items online for the user based on his/her verbal instructions into their mobile device. The goal is to build a mobile Web-app that is more useful and increases the convenience of the user, while at the same time protecting his or her personal data in a transparent and seamless manner based on his/her preferences and the context of the situation. Contextual processing is mandatory for effective privacy protection, as well as in other data-related applications.

We foresee that such a breakthrough is the primary challenge in building “smart” apps and, as such, would serve as a platform technology for expansion into many areas. Building a virtual machine that is able to act based on its ability to “view” its surroundings as a series of interactional contextual slices of the world is the true “killer-app” of the 21st century. Our strategy, accordingly, is not to focus on privacy alone, since such apps will be limited to a small market segment of “privacy-fundamentalists;” rather, we will embed privacy into smart Web-apps. We suggest that is a far more effective marketing strategy in rolling out privacy-protecting devices to consumers. By incorporating the principles of SmartData into Web-apps such that they now incorporate convenience/utility *plus* privacy protection, the likelihood of adoption across a wide customer base will be enhanced.

2.5 SmartData’s Structure: Zooming in

The security infrastructure for SmartData will consist of first stripping a user’s personal ID from the body of his/her data. The anonymized data is then segmented, encrypted, and placed in digital “lock-boxes.” Each lock-box will have a non-personal

description of its segmented data as an attached meta-tag; for example, MRI imaging data, address information, and so forth. These lock-boxes will be stored in the cloud, with the agent storing an encrypted pointer to the appropriate location—analogous to an anonymous database structure. The SmartData agent will store the pointers, encryption keys, and the meta-tags. Access to this information will be controlled by the agent and “designated individuals” in the event of emergencies. The challenge will be to securely store keys within the agent’s structure such that it has access to them when necessary, analogous to a person remembering his/her passcode to access a secured building.

Once SmartData is developed (evolved and programmed), the objective is to have a binary string specify its architecture, both neural and somatic.¹ The binary string will contain both the code (weights, etc.) for the neural architecture which will embody the “smartness” and the encryption keys and pointers. It will also specify the pre-programmed modules specifying the language lexicon and transaction instructions as described below. As such, there will be no different pieces of the binary string; it will not contain an identifiable segment representing sensitive data. The *entire* string will specify the neural and somatic architecture of the agent. This binary string or code, when downloaded into a hardware device that is reconfigurable will restructure the device into the architecture specified by the binary code, and “embody” a unique agent which will serve as the cyber-surrogate for the data subjects. Banks of these reconfigurable devices will be located in the cloud and accessed by users when required to perform Web-services. A hypothetical analogy is that the binary string stored in the cloud contains the type and location of all the atoms in a person’s body which when downloaded into a cloning machine will reconstruct a clone of the person. That clone, which contains sensitive information stored within the connections of its neuronal network in the brain, can now act as the person’s surrogate to perform services in a privacy-protective manner. Similarly, SmartData will incorporate the pointers to the anonymized personal data (stored in the cloud) into its neural architecture. It can therefore serve as a surrogate of its owner to either protect or release its data, based on its owner’s instructions and the context of the request. In other words, the data becomes “smart” by virtue of its being inextricably linked to, and secured within the intelligent agent.

2.6 The Background to the Approach

The question before us is: How do we build agents that “live” in cyberspace and are smart enough to protect the personal data of a user when carrying out a Web-service on their behalf? We look to natural evolution as providing us with a general template for our approach.

¹ Somatic refers to a body. We will define below what a body is in cyberspace.

Nature did not start off by evolving smart minds; it evolved bodies and associated actions first—from bacteria to reptiles, mammals and so forth. In other words, it evolved morphology, and “minds” followed. Nature, in selecting the morphology of the body, constrains actions into a finite number of degrees of freedom. Our survival, both as a species and individuals, is directly related to the number of different ways that we can respond to situations that we encounter in our world. Expansion of our degrees of freedom was “selected for”² in nature, initially by evolving appendages and articulated joints, with increases in the cognitive capacity to optimize the use of these physical appendages. With increased cognitive capacity, primates, and especially humans, started to design, make, and use tools to further increase their degrees of freedom when interacting with the world. Tools extended our minds by giving us new skills which, in turn, further augmented our degrees of freedom and increasingly opened up the world to us. The reason that we are bringing this up here is because computer languages and programs are artifacts or tools, which as you will see below, gives us a novel approach in amalgamating bottom-up evolutionary computation with top-down machine learning techniques in building SmartData agents.

It is obvious that humans playing tennis or birds flying is a direct consequence of their morphology. But what may not be so obvious is that the body can be regarded as a set of potential actions which when smoothly sequenced together form behaviours. In other words, behaviour can be viewed as the selection of a sequence of elementary actions from the total set of actions embodied in the morphology. When I reach for a cup of coffee, take a sip, and put it back, that movement is the orchestration of a sequence of muscle actions in various parts of my anatomy. Each action is elementary on its own, but when strung together achieves my objective of drinking coffee. In effect, a body is nature’s way of acting in its “world” through the set of evolved actions as represented by its morphology. What this says is that nature was not evolving bodies per se, it was evolving actions. And because actions are synonymous with work, which is a function of force and force is a function of mass, then actions have to be embodied within vehicles of “mass” in our “real gravity-influenced world.”

However, it is not a mere set of actions that define a body; it is a set of actions which, in the case of biological agents, has been evolved specifically to satisfy the agent’s needs in its quest for survival. A body is defined always from the agent’s perspective. This says that a body is what a body does. And what does a body do? It acts to satisfy needs. A body that cannot act in this way is not a body; it is a non-autonomous machine or a corpse. Therefore, self-serving actions define a body and different sets of potential actions define different species. When we say that an agent is embodied, we are really saying that it has the inbuilt functionality (a body) capable of performing autonomous actions to satisfy its needs.

² Actually, nature did not select for, it selected against, by virtue of decreased offspring relative to other individuals.

Consequent to this view, nature spent the major part of our evolutionary history evolving biological agents, from snails to humans, with different sets of actions. The cognitive capacity of each species of agents was a direct function of the morphology or more accurately, set of evolved actions. The greater the degrees of freedom represented in the morphology of a species, the greater the potential cognitive capacity that could be achieved given an appropriate and demanding environment. The rationale behind this, which was elucidated by Llinas [4], is that there is a system or network(s) in the brain that evolved to select the elementary actions (he called them fixed-action patterns) constituting any behaviour. This neuronal network can be modeled by a dynamical system where its trajectory in a basin of attraction constitutes a sequence of neuronal states whose outputs to the cranial nerves and spinal cord selected the elementary actions making up a gross behaviour or movement. As morphology evolved and expanded the degrees of freedom of the body (and the set of potential elementary actions), the network, by necessity, also had to evolve and become more complex in order to incorporate its expanded repertoire.³ Llinas postulates, however, that the evolved complexity resulting from expanding the degrees of freedom of the organism was “seconded” by the mental system, so to speak. He states: “That which we call thinking is the evolutionary internalization of movement.” In other words, the dynamical network selecting elementary actions is isomorphic across domains and, as such, is adept at selecting both “mental” as well as physical actions.

This proposition provides us with a strategy that may shorten the developmental process of SmartData. SmartData will be operating in cyberspace and, as such, there is no gravity and therefore “mass-type” actions are not required as they are in our “real” world. We posit that in cyberspace, the parallel to a physical body’s set of actions is a set of transaction instructions or codes that performs actions on the Web that satisfies the needs of SmartData. This set of Web-related instructions comprises SmartData’s body. This implies, based on our view of embodiment, that SmartData has “needs” (which we will address below). Therefore, behaviour by a SmartData agent is the selection of a sequence of transaction instructions which accomplishes its objective. And if SmartData is the surrogate of the data-subject, those objectives are the data subjects’ objectives. Accordingly, if SmartData is to be an autonomous agent, then the needs of the data subject must somehow be placed into its cognitive structure.

If the set of transaction instructions (which constitutes SmartData’s body) are such that different sequences selected from them exhaust all possible behaviours to achieve potential objectives required in a domain (a domain such as ordering books online in a privacy-protective manner), then SmartData’s “body” can operate effectively in its domain/environment. However, since in a dynamic environment, it may be impossible to forecast all of the possible behaviours required to satisfy all

³ We say “by necessity” because, if for example, an agent evolved an additional limb which it did not use, it would only serve as an additional energy drain (weight, metabolism, etc.) without a concomitant benefit. Therefore, it would be selected against.

potential future objectives, or in our case, requests to perform services or release personal data, it is important that the set of instructions be generalizable. This means that each transaction instruction must be elementary or simple enough so that stringing together, or sequencing, a number of them can generate any potential behaviour required for a particular domain—in our case the Web-domain.

This method affords us with a potential shortcut since it precludes the necessity to evolve a body or set of actions. Instead of evolving a body, we judiciously choose (hand-code) a set of instructions which when sequenced into a subroutine will handle all potential actions required in a Web-domain. A related example would be opcode mnemonics instructions in Assembler language where sequences of these instructions (actions) have satisfied all programming objective so far conceived. Forty instructions, for example, will generate well over one trillion sequences. Though a large percentage would not be useful, even a few percent of a trillion instructions that are, would comprise a large number of behaviours.

Accordingly, our prime task is to evolve a dynamical recurrent network which, based on input requests, undergoes transitions in its state (its trajectory) which serves to select the appropriate sequences of instructions (which are hand-coded beforehand). As an example, if we assume that the neural network has eight outputs which are binary, then with an eight-bit output, up to 256 instructions may be selected. As the state of the network travels along its trajectory, it will select a sequence of instructions based on the value of the eight-bit state. Analogous to symbolic dynamics, we are coarse coding the state-space of a dynamical system into a finite set of partitions or subspaces and assigning an eight-bit symbol to each. Whenever the system crosses a partition, the assigned eight-bit symbol is produced. In this way, trajectories can be represented as sequences or strings of symbols which select Web-based instructions. Such systems have a built-in redundancy since there may be more than one state between partitions. We believe, however, that the redundancy inherent in such networks will be valuable in that more than one state or symbol can select the same instructions. This may reduce the probability of getting trapped in a local minimum within search space.

In a way, nature already uses an analogous method when it evolved language, for example. Language is a sequence of speech sounds that is coded into a set of n distinct, concatenated phonemes which can be viewed as symbols. Therefore, when listening, the sensory input is a sequence of n symbols. Conversely, when speaking, the internal representations of these symbols trigger motor actions. In English, any sequence of vocal actions is selected from a set of about 40 phonemes. One can view all vocal output as the result of the selection, sequencing and timing of around 40 instructions or mini-subroutines. The instructions are the phonemes and comprise the set of actions. However, nature not only had to evolve the neural networks to select the sequence of phonemes, it also had to evolve (biologically and culturally) the actual phonemes, something that we would rather not have to do.

2.7 SmartData as a Mobile App for Online Web Services

We bring in language at this point because the utility of mobile apps on the Web will be significantly enhanced if they can operate accurately based upon voice requests.

Language evolved consequent to a long series of biological as well as cultural adaptations. Agents were situated within a world with other agents in order for these adaptations to be selected. To emulate nature's accomplishment of evolving both the anatomical structure (biological evolution or phylogeny), and the phonetic lexicon and grammar (cultural evolution or glossogeny) requires a community of agents evolving in a world—especially to evolve a shared lexicon and grammar. Furthermore, the learning stage (ontogeny) also depends on a community of agents within which to learn and practice language. However, language learning can only take place when both the anatomical “equipment” and the lexicon/grammar already exist; the former in the agent, and the latter in “society.” We do not yet have the capability to accomplish this for cyber-agents. Our proposed method, therefore, is to bypass the need to evolve an “anatomy” in agents and a lexicon/grammar in a “community” of agents. We will instead “construct” the anatomy (the microphone and speaker built into the mobile device), and create a finite-symbol formal language for mobile book ordering. The lexicon/grammar of the formal language will be hand-coded in a manner similar to our set of transaction instructions for the Web-domain. Initially, we will construct a formal language with a limited lexicon with morphemes or terms related to book ordering, and in later phases of the development, try to expand the formal language more toward a natural language.

In general, learning a language can be viewed as analogous to unsupervised learning of a discrete stochastic process—as an example, a black box with n lights representing the n terms in our lexicon. Concealed within the black box is the generative process underlying the language which, in this case, causes different lights to sequentially flash-on at discrete intervals of time, similar to a sequence of terms in a verbal request. The goal of learning is to infer the underlying nature of the stochastic process within the black box. However, that is only part of the task. The lights on the black box are already separated or “parsed” and can be analyzed as distinct events. Initially, that is not the case with spoken language. In listening to an unknown or foreign language it is sometimes difficult to differentiate the sounds into distinct terms. This is a learning process in itself which, although in natural language is part and parcel of the entire learning process, can be subdivided into a separate stage. This is our strategy with SmartData.

The first stage of differentiating the terms will be accomplished using a supervised learning process. In the SmartData mobile app setup, the sensor transducer will be a microphone with an analog output which will be fed into a language term “parsing” network whose purpose will be to parse the stream of speech from the user into terms and select (analogous to selecting transaction subroutines) the appropriate sequence of hand-coded lexicon codes representing the terms in the verbal instructions. Selection of the terms will be accomplished by learning/memorizing the lexicon codes associated with each term of the

sequence.⁴ The “parsing neural network” will be evolved to select the correct sequence of lexicon codes, L_1, \dots, L_n , matching the actual verbal instructions, also selected from the lexicon. The training input data, X , will be a set of different verbal instruction, using the English terms in the lexicon, spoken by different people. The training output data, Y , will be the predefined lexicon codes L_1, \dots, L_n . This, in a way, is analogous to the neural codes from the prefrontal gyri in our brains which extracts (filters) the sequence and timing of phonemes and syllables from the analog signals into the ear (via the cochlear nucleus). Since our formal language has a fixed lexicon that is not arbitrary or dynamic, as with natural languages, this may preclude the necessity to evolve these agents in a community, which hopefully, will simplify our task.

The sequence of codes from the predefined lexicon (which will represent a user’s verbal instructions), together with the correct sequence of transaction instruction codes (which will carry out the appropriate Web-related actions), will comprise the training data set for the behavioural or transactional system which we term the Landscape Generation Network (“LGN”). However, the LGN will comprise two sub-networks; a context network whose inputs will be the sequence of lexicon codes and whose function will be to generate a control parameter which will be the input to a second transactional sub-network. This second transactional sub-network will output a sequence of codes selecting the pre-defined transactional instructions. The sequence of transactional codes along with information from the Web page regarding status of the transaction will also be fed back to the context network. The objective is to modify the context control parameter based on the transaction instruction outputs and the response from the Web page to those outputs (e.g. error messages). The context control parameter chooses (in a dynamical systems sense) specific basins of attraction or landscapes in the transactional sub-network. The trajectory of the state on the landscape will comprise the codes that select the sequence of instructions. In effect, we will be attempting to evolve a set of partitions which will course code the state space and provide us with code redundancy as with the lexicon codes. All three networks—language parsing, context and transactional—will be evolved, and the fitness function will select for correct sequences of transaction instructions based on input verbal instructions, with the objective of minimizing out-of-sample error.

Part of our task will be to determine to what extent we can evolve these landscapes networks (which are dynamical recurrent neural nets with the genes controlling the number of neurons in a network and the Hebbian rules for the synapses) to be compact in the mathematical sense. This is important since compactness infers generalizability. In other words, if we train our agent to respond correctly to a partial set of verbal instructions (training set) with a low in-sample error rate (and where the number of training examples is much greater than the number of synaptic weights time the logarithm of the number of neurons in the network), then we have a high probability that it will also respond correctly to

⁴In this case, it is more accurate to label it a supervised learning/memorizing process. The parsing is a learning process while the attachment of a lexicon code is memorization.

inputs not seen before, that is, the out-of-sample error will be bounded by the in-sample error plus some delta factor. This will, in part, depend on the VC dimension [5] and the homogeneity of the domain. We posit that in biological agents, nature selected and “categorized” a set of objects or events as belonging to a particular domain using interactive context as a parameter to somehow “alter” the effective input space and thereby “structure” seemingly inhomogeneous objects/events as “related” such that a compact description could be evolved. This will be part of our task in this research—to see how we can utilize interactional context to, in effect, “homogenize” a domain.

In summary, we will essentially evolve agents to parse verbal instructions by evolving landscapes to select the proper sequence of lexicon codes based on the input instructions. These lexicon codes will then form the inputs into a context landscape sub-network which will output a control parameter as an input to a transactional landscape sub-network in which the sequence of neural states (trajectory) selects appropriate transaction instructions for the Web to satisfy the verbal instruction. This methodology we believe, will generalize to other domains. In another domain, we will have to either add a new lexicon and transaction subroutines specific to the new domain to the existing set, or substitute the lexicon and transaction sets with new domain sets.

2.8 The Significance of Context

Everything that we do is a function of context. “Thus, much of the meaning of a represented piece of information derives from the context in which the information is encoded and decoded” [3]. In fact, to an intelligent agent, the world is a temporal series of contexts and these contexts solicit actions by bringing the appropriate landscape to the foreground. And those actions can be thoughts, words or deeds. To understand this statement we have to unpack some of the concepts and we will do that within the context of privacy. Ideally, when there is a data request, the data-subject can make a decision to either release the information or not, or to release with certain restrictions. But to make that decision he/she may consider a number of factors, such as: the purpose and use of the data e.g., applying for a passport or trying to join a social network; the identity of the requestor; authorization and authentication of the requestor; the type of data requested; and so forth. In a nutshell, there is a request for a specific type of personal data, and there is a context to the request. We are defining context as *related to* those factors listed above.

What we mean by “related to” is that the significance of any of these factors in the decision to release personal data is *relative* to the overall needs,⁵ preferences,

⁵ Henceforth, I will refer to “overall needs, preferences, and previous actions of the data subject together with the feedback from the environment relative to the specific action” as embodied needs.

and previous actions of the data subject together with the feedback from the environment relative to the specific action.⁶ If the purpose of the personal data is to obtain a passport, for example, and I really need a passport, then the significance or “purpose vs. need” score for that factor will be high and I will release data, all other things being satisfactory on that list. If I do not need a passport, on the other hand, the significance will be low and I won’t release the personal data. Context scores the significance of a factor using, as a benchmark, the embodied needs of the agent. But that benchmark of embodied needs, itself is dynamic.

So this is one of the challenges—to somehow initially infuse SmartData with a proxy of the data subject’s embodied needs, so that context has real and accurate significance. We then require SmartData to modify those needs based on its experience, within the data subject’s constraints.

2.9 Conclusion

There is yet another factor to consider which will bear on the eventual success of SmartData. We know that the practices and procedures involved in safeguarding privacy and security on the Web are derived from individuals’ concerns and solutions in the “real” world. These concerns and solutions are themselves derived from the social and cultural environments in which we live. Therefore, if an agent in cyberspace is to function autonomously and effectively it must first “understand” the specific social and cultural environment of humans within the domain in which it will operate. An “electronic health-care agent” may not need to understand the social and cultural environment of professional basketball; it must, however, understand the environment within the domain of health-care. Hence, at some point in our development, we suggest that agents must be evolved within a simulated virtual world that presents the relevant attributes of the domain in which it will operate such that the proper cognitive characteristics will be selected. One cannot just “program in” or encode relevant contexts as has been the practice in standard AI when applied to narrow and static domains. Furthermore, SmartData agents must at some point in their evolutionary cycle inhabit a world with other agents in order to allow for inter-subjective cooperation and competition which, as has been demonstrated in our evolution, gives rise to particular social practices and cultures. These are formidable long-term challenges, requiring considerable innovation—a great deal of scientific ground breaking will have to occur.

⁶ By actions, I do not mean physical actions alone but also “mental actions” such as thought or listening to someone speak. In the case of speech, what I have just heard in the immediate past serves as the benchmark to score the significance of what I am hearing now and therefore alters the landscape such that what I hear in the future will be contextually biased.

However, these very challenges position it among the most exciting research one could think of undertaking! And it comes with enormous payoffs—privacy and civil liberties for one, but also the myriad of innovative spin offs of processing information contextually in a manner that is natural for biological agents.

References

1. Cavoukian, A. (2006). *Creation of a Global Privacy Standard* (pp. 4).
2. Cavoukian, A. (2012). *Why are We Here Today? Privacy and the Promise of SmartData*. IPSI SmartData International Symposium, Toronto, Ontario.
3. Lenat, D. (1998). *The dimensions of context-space*. CYCorp Report.
4. Llinas (2001), R.R., *I of the Vortex: From neuron to self*, MIT Press Cambridge Mass.
5. Vapnik, V.N. and Chervonenkis, A. Y, (1971) *On the uniform convergence of relative frequencies of events to their probabilities*, Theory of Probability and its Applications, 16, pp. 264–280.

SmartData

Privacy Meets Evolutionary Robotics

Harvey, I.; Cavoukian, A.; Tomko, G.; Borrett, D.; Kwan, H.; Hatzinakos, D. (Eds.)

2013, X, 214 p., Hardcover

ISBN: 978-1-4614-6408-2