

INFORMATION ETHICS GROUP
Oxford University and University of Bari

Internet Ethics: the Constructionist Values of Homo Poieticus

by

Luciano Floridi and J.W. Sanders

luciano.floridi@philosophy.ox.ac.uk

jeff@comlab.ox.ac.uk



IEG – Research Report 25.02.03

<http://web.comlab.ox.ac.uk/oucl/research/areas/ieg>

Abstract

In this chapter, we argue that the web is a poietically-enabling environment, which both enhances and requires the development of a “constructionist ethics”. We begin by explaining the appropriate concept of “constructionist ethics”, and analysing virtue ethics as the primary example. We then show why CyberEthics (or Computer Ethics, as it is also called) cannot be based on virtue ethics, yet needs to retain a constructionist approach. After providing evidence for significant poietic uses of the web, we argue that ethical constructionism is not only facilitated by the web, but is also what the web requires as an ethics of the digital environment. In conclusion, we relate the present discussion to standard positions in CyberEthics and to a broader project for Information Ethics.

Keywords

Computer Ethics, CyberEthics, constructionism, Information Ethics, poiesis, Virtue Ethics, web.

Note

Forthcoming in R. CAVALIER (ed.), *The Impact of the Internet on Our Moral Lives* (New York: SUNY, 2003).

©2003. L. FLORIDI - J.W. SANDERS, *Internet Ethics: the Constructionist Values of Homo Poieticus*, IEG Research Report 25.02.03, digital editing by G. M. Greco, Information Ethics Group, Oxford University – University of Bari, <http://web.comlab.ox.ac.uk/oucl/research/areas/ieg>.

“Go create”

Sony Advertisement, 2002

1. Introduction

Ethical issues are often discussed in terms of putative resolutions of hypothetical situations, such as “what should one do on finding a wallet in the lavatory of a restaurant?”. Research and educational purposes may promote increasingly dramatic scenarios (sometimes reaching unrealistic excesses¹), with available courses of action more polarised and less easily identifiable as right or wrong. But the general approach remains substantially the same: the agent is confronted by a moral dilemma and asked to make a principled decision by choosing from a menu of alternatives. Moral action is triggered by a situation.

In “situated action ethics” (to borrow an expression from AI), such moral dilemma may give the false impression that the ethical discourse concerns primarily *a posteriori* reactions to problematic situations in which the agent unwillingly and unexpectedly finds herself. The agent is treated as a world user, a game player, a consumer of moral goods and evils, a browser,² a guest, or a customer who reacts to pre-established and largely unmodifiable conditions, scenarios and choices. Only two temporal modes count: present and future. The past seems irrelevant (“how did the agent found herself in such predicament?”), unless the approach is further expanded by a casuistry analysis.

¹ See, for example, “the trolley problem” (Foot [1967] and Thomson [1976]; for a very entertaining parody do not miss “the revised trolley problem” in Patton [1988]). On “George’s job” and “Jim and the Indians” see Smart and Williams [1973]. Contrary to the trolley problem, the last two cases are meant to provide counterexamples against purely consequentialist positions.

² For an entirely “situation based ethics” approach to the Internet see for example Dreyfus [2001]. Dreyfus seems to ignore entirely any constructionist issue. His “anthropology” includes only single web users browsing the net.

But ethics is not only a question of dealing morally well with a given world. It is also a question of constructing the world, improving its nature and shaping its development in the right way. This *proactive*³ approach treats the agent as a world owner, a game designer or referee, a producer of moral goods and evils, a provider, a host, or a creator. The agent is supposed to be able to plan and initiate action responsibly, in anticipation of future events, in order to (try to) control their course by making something happen, or by preventing something from happening rather than waiting to respond (*react*) to a situation, once something has happened, or merely hoping that something positive will happen.

There are significant differences between reactive and proactive approaches. There is no space to explore them here, but one may mention, as a simple example, the moral responsibilities of a webmaster as opposed to those of a user of a web site. Yet, differences should not be confused with incompatibilities. A mature moral agent is commonly expected to be both a morally good user and a morally good producer of the environment in which she operates, not least because situated action ethics can be confronted by lose-lose situations, in which all options may turn out to be morally unpleasant and every choice may amount to failure. A proactive approach may help to avoid unrecoverable situations. It certainly reduces the agent's reliance on moral luck. As a result, a large part of an ethical education consists in acquiring the kinds of traits, values and intellectual skills that may enable the agent to switch successfully between a reactive and a proactive approach to the world.

All this is acknowledged by many ethical systems, albeit with different vocabulary, emphasis, and levels of explicitness. Some more conservative ethical theories prefer to concentrate on the reactive nature of the agent's behaviour. For example,

³ We use the term "proactive" technically, to qualify policies, agents, processes or strategies that (a) implement effective action, in anticipation of expected problems, difficulties or needs, in order to control and prevent them, at least partially, rather than merely reacting to them as they occur (in this sense an ethically proactive approach can be compared to preventive medicine, which is concerned with reducing the incidence of disease by modifying environmental or behavioral factors that are causally related to illness); or that (b) actively initiate good changes, promoting rather than merely waiting for something positive to happen.

deontology embeds a reactive bias insofar as it supports duties on-demand. Another good example is the moral code implicit in the Ten Commandments, which is less proactive than that promoted in the New Testament. On a more secular level, the two versions of Asimov's laws of robotics provide a simple case of evolution. The 1940 version is more reactive than the 1985 version, whose new zeroth law includes a substantially proactive requirement: "A robot may not injure humanity, or, through inaction, allow humanity to come to harm" (see Clarke [1993-4] for a full analysis and further references).

Ethical theories that adopt a more proactive approach can be defined as *constructionist*. They are the ones that interest us here. One of the best examples of constructionist ethics is virtue ethics. The analysis of its scope and limits will introduce our discussion of a constructionist approach to CyberEthics.

2. The scope and limits of virtue ethics as constructionist ethics

According to virtue ethics, an individual's principal ethical aim is to live the good life by becoming a certain kind of person. The constructionist stance is expressed by the desire to mould oneself. The goal is achieved by implementing or improving some characteristics, while eradicating or controlling others. The stance itself is presupposed: it is simply assumed as uncontroversial that one does wish to live the good life by becoming the best person one can. Some degree of personal malleability and capacity to choose critically provide further background preconditions. The key question "what kind of person should I be?" is (rightly, in our view) considered to be a reasonable and justified question. It grounds the question "what kind of life should I lead?" and immediately translates into "what kind of character should I construct? What kind of virtues should I develop? What sort of vices should I avoid or eradicate?". It is implicit that each agent strives to achieve that aim *as an individual*, with only incidental regard to the enveloping community.

Different brands of virtue ethics disagree on the specific virtues and values identifying a person as morally good. The disagreement, say between Aristotle, Paul of Tarsus and Nietzsche, can be dramatic, not least because it is ultimately ontological, in that it regards the kind of entity that a human being should strive to become. In

prototyping jargon, theories may disagree on the abstract specification of the model, not just on implementation details.

Despite their divergences, all brands of virtue ethics share the same subject-oriented kernel. This is not to say that they are all subjectivist but rather, more precisely, that they are all concerned exclusively with the proper *construction* of the moral subject, be that a self-imposed task or an educational goal of a second party, like parents, teachers or society in general. To adopt a technical expression, virtue ethics is intrinsically *egopoietic*. Its *sociopoietic* nature is merely a by-product, in the following sense. Egopoietic practices that lead to the ethical construction of the subject inevitably interact with, and influence, the ethical construction of the community inhabited by the subject. So, when the subjective microcosm and the socio-political macrocosm differ in scale but essentially not in nature or complexity, as one may assume in the idealised case of the Greek *polis*, egopoiesis can scale up to the role of general ethics and even political philosophy. Plato's *Republic* is an excellent example. Plato finds it unproblematic to move seamlessly between the construction of the ideal self and the construction of the ideal city-state. But so does the Mafia, whose code of conduct and "virtuous ethics" for the individual is based on the view that "the family" is its members.

Egopoiesis and sociopoiesis are interderivable only in sufficiently *simple* and *closed* societies, in which significant communal behaviour is ultimately derivable from that of its constituent individuals. It is hard to specify "sufficiently" precisely, but some light can be cast here by trying to clarify what "simple" and "closed" mean.

On the one hand, "simple" refers to the "vertical" growth of a society, that is to its degree of *autonomy*. A society is no longer sufficiently simple, but qualifies as increasingly *complex*, when some of the major new variables that govern its development are internal forces, emerging holistically from the actions and decisions of its members, forces like unemployment or price inflation, for example, which are beyond the control of single human agents.

On the other hand, the threshold between a closed and an open society (no reference to Popper here) is to be identified in the level and relevance of interconnections and interactions between the society in question and other similar macro-agents. A sufficiently open society is one in which some of the major new variables that govern its

development are external forces influencing it from without. Therefore, “open” and “closed” indicate the relative degree to which interaction determines evolution. This is the “horizontal” growth of a society.

Societies exhibit a continuum of stages, with simple and closed societies at one end of the continuum and, at the other end, societies open and complex enough to sustain autonomous behaviour and demonstrate *emergent*⁴ characteristics. As they evolve, societies may progress along the continuum. At some point, whilst immediate and personal interactions among all its members are still significant, in practice systemic forces may supervene, profoundly influencing the life of each individual. Such open and complex societies inherit from their constituent individuals *autonomy* and *interactivity* and, at a certain level of evolution, they become *adaptive*. They thus form (*artificial*) *agents* by virtue of those three properties (on the ethics of artificial agents see Floridi and Sanders [2001b]). In such societies, sociopoiesis is no longer reducible to egopoiesis alone. This is the fundamental limit of virtue ethics. In autonomous, interactive and adaptive societies, virtue ethics positions acquire an individualistic value, previously inconceivable, and may result in moral escapism. The individual still cares about her own ethical construction and, at most, the construction of the community with which she is more closely involved, like the family, but the rest of the world falls beyond the horizon of her moral concern.

All this was true during the last centuries of the Roman Empire, for example, and applies equally well in our new era of globalisation. Phrasing the point in terms of situated action ethics, new problematic hypothetical situations arise from emergent phenomena. Examples include issues of disarmament, the ozone level, pollution, famine and the digital divide. The difficulty becomes apparent in all its pressing urgency as the

⁴ Communal behaviours that are not immediately or directly so explicable are called emergent. Perhaps the simplest examples come from artificial communities. In Conway’s Game of Life, for example, the behaviour of an individual is determined by the states of its immediate neighbours. Stable, periodic or otherwise interesting behaviour (e.g. gliders, which retain their collective state but glide across the digital landscape) of subcommunities consisting of several individuals provide examples of emergent behaviour. In our own, real, global society, monetary inflation, unemployment and such phenomena whose dynamics are determined by feedback of data from subcommunities provide examples.

individual agent tries to reason using “local” ethical principles to tackle a problem with “global”, ethical features and consequences.

3. Why CyberEthics cannot be based on virtue ethics

We are now in a position to distinguish between two phenomena often confused in the literature on CyberEthics: the renewed popularity of virtue ethics (a) in our society (see Slote [2000] for a sympathetic overview) and (b) in cyberspace (Coleman [1999], [2001]; Grodzinsky [2001]).

In case (a), one is confronted by a context in which an individualistic culture facilitates practically, but does not justify theoretically, the return to a subject-oriented ethics. One should still properly object that (i) the kind of egopoiesis promoted by virtue ethics cannot (indeed, was not meant to) scale to very complex and open social contexts; and (ii) virtue ethics presupposes a philosophical anthropology (a theory of what it means to be fully human) that, in a sufficiently evolved social context, cannot be left embedded but that, once it is made fully explicit, requires an ethical justification to become acceptable precisely as a *morally good* anthropology, and hence as ethically preferable.

In case (b), phenomena like the great popularity of “virtual communities” (see section 4.5), which arguably represent the digital re-incarnation of the polis, mean that people naturally tend to concentrate on the ethical construction of their “personae” as, at the same time, a contribution to the construction of the agent’s self and a substantial contribution to the construction of the local cyber-community, which is largely characterised by the members constituting it and inhabiting it. In this simple and closed context, an egopoietic approach is indeed fruitful, precisely for the same reasons it was in the polis. One is justified in arguing that virtue ethics may be all that is needed for the ethical well-being of the whole community.

The two trends (a) and (b) have merged and currently interact in the information society, but they are better understood separately, lest one should mistakenly argue that because virtue ethics can work in small cyber-communities (comparable to local area networks) and it is popular “IRL” or “OT” (in real life or out there) it is also all that CyberEthics needs as a theoretical foundation. The opposite is true. Because virtue ethics remains limited by its subject-oriented approach and its philosophical anthropology, it cannot provide, by itself, a satisfactory ethics for a globalised world in general and for the information society in particular. If misapplied, it fosters ethical individualism, as the agent is more likely to mind only her own self-construction. If it is uncritically adopted, it

can be intolerant, since agents and theorists may forget the culturally over-determined nature of their foundationalist anthropologies, which often have religious roots. If it fosters tolerance, it may still spread relativism because any self-construction becomes acceptable, as long as it takes place in the enclave of one's own private sphere, culture and cyber-niche, without bothering any neighbour.

The inadequacy of virtue ethics is of course historical. The theory has aged well, but it can provide, at most, a local sociopoietic approach as a mere extension of its genuine vocation: egopoiesis. It intrinsically lacks the resources to go beyond the construction of the individual and the indirect role this may play in shaping her local community. Theoretically, however, the limits of virtue ethics should not lead to an overall rejection of any constructionist approach. On the contrary, the fundamentally constructionist lesson taught by virtue ethics (one of the features that make virtue ethics appealing in the first place) is more important than ever before.

In a global information society,⁵ the individual agent (often a *multi-agent system*⁶) is like a demiurge. Her ontic powers can be variously exercised (in terms of control, creation or modelling) over herself (e.g. genetically, physiologically, neurologically and narratively), over human society (e.g. culturally, politically, socially and economically) and over natural or artificial environments (e.g. physically and informationally). Such an increasingly powerful agent has corresponding moral duties and responsibilities to oversee not only the development of her own character and habits but also the well-being of each of her spheres of influence. Clearly, a constructionist ethics should be retained and reinforced. The mistake (developing CyberEthics in terms of virtue ethics) lies not in the stress put on constructionism *per se*, but in the direction in which constructionism is presupposed to develop: namely only towards the individual source of the moral action (building the character of a human agent) instead of the receiver of the moral action as well, that is towards the patient, the object and more generally the environment affected

⁵ On the history of the development of the global information society see Mattelart [2001].

⁶ A multi-agent system (MAS) is a conglomeration of interacting components, known as agents, capable of cooperating to solve problems that typically are beyond the individual capabilities or knowledge of each agent. Thus a MAS exhibits a greater system-level behaviour than its constituting agents (Huhns and Singh [1998]).

by the action. The kind of ethical constructionism needed today goes well beyond the education of the self and the political engineering of the simple and closed cyberpolis. It must also address the urgent and pressing question concerning the kind of global realities that are being built.⁷ This means decoupling constructionism from subjectivism and re-orienting it to the object, applying it *also* to society and the environment, the receivers of the agent's actions.⁸

The term "ecopoiesis" refers to the morally-informed construction of the environment based on this object- or ecologically-oriented perspective. To move from individual virtues to global values an *ecopoietic* approach is needed that recognises the agent's responsibilities towards the environment (including present and future inhabitants) as its enlightened, creator steward or supervisor, not just as its virtuous user and consumer.

An ecopoietic ethics, like any form of constructionism, raises a fundamental ontological concern. Moral luck aside, the chances of constructing an ethically good *x* increase the better one knows what an ethically good *x* is, and vice versa. Constructionism depends on a (satisfactory epistemic access to, or understanding of, the) relevant ontology. In the context of digital environments, an ecopoietic ethics presupposes a substantial answer to the foundationalist question "what is the essential nature of information, computers and the internet?". If virtue ethics presupposes a philosophical anthropology, an ecopoietic ethics seems to require a *Philosophy of Information* ([Floridi, 1999], [2002] and [2003a]). In the rest of this article, we shall not pursue this ontological foundation of constructionism. Instead of looking at the theoretical roots of constructionism, we shall concentrate on its branches, and seek to clarify the connection between CyberEthics and constructionism by showing how the

⁷ We have addressed the issue of the construction of ethical artificial agents in Floridi and Sanders [2001b].

⁸ In Floridi [1999] and [forthcoming], we have argued that this is in line with the development of contemporary ethics, which has registered a general shift from the centrality of the agent in standard macroethics such as virtue ethics, deontology, consequentialism and contractualism, to the centrality of the patient in non-standard macroethics such as environmental ethics, bioethics and medical ethics.

latter emerges from the web⁹ and how the web can benefit from a constructionist approach.

4. Poiesis on the web

The web is changing patterns of moral behaviour in many ways, with important repercussions on the development of the ethical discourse. Instances of situated action ethics, primarily with negative consequences, have attracted a large variety of detailed analyses, and account for most of the literature in CyberEthics (see for example Spinello and Tavani [2001] and other chapters in the present volume). The web, however, is not only a source of moral dilemma. As a new social space and digital environment, it has also greatly enhanced the possibility of developing egopoietic, sociopoietic and ecopoietic projects. It has thus contributed to the emergence of a constructionist ethics as a macroscopic phenomenon. In this section we shall consider a range of indicative examples, which well illustrate the ethics of constructionism.

4.1 Interfaces

Choosing and modelling one's own interface to the digital world represents a first, indicative example of the kind of constructionism promoted by the web. A user's most immediate interactions with the web lie with an interface, whose features therefore influence her view. By tradition, a well-designed interface offers its user a convenient *mental model* for the actions it supports. For instance, one design principle states that, if an action has different effects in different situations, the prevailing *mode* that determines the effect should be intuitively clear to the user. Typical mental models in this context are the "desktop", "folder" and "filing cabinet". As an example of the model clarifying the mode, by adopting the mental model of the text file as a folder the user is able to appreciate that the depression of a key has different effects when a text file is open or when it is closed. On the other hand, that model is limited because it does not address why the user needs periodically to "save" the results of editing the file.

⁹ For current purposes no distinction is drawn between the internet and the web.

Laurel (Laurel [1991]) has proposed an alternative view of interfaces as theatre, following Aristotle's six elements of drama. In order of increasingly abstract material cause (that one of Aristotle's four causes, operating during the process of creation, which reflects the fabric from which a thing is made), together with their interpretation in human-computer activity, they are (adapted from Laurel [1991], Table 2.1):

MATERIAL CAUSE	INTERFACE ACTIVITY
Spectacle/enactment	all sensory components of the action represented
pattern/melody	the pleasurable perception of pattern in the sensory phenomena
Language/diction	the selection and arrangement of signs, used semiotically
Thought/reasoning	the inferred internal processes leading to choice, of both human and computer
Character/agency	the bundles of predispositions and traits, of both human and computer
plot/action	the whole action; a collaboration between system and user.

This approach places emphasis on designing the action (to be engaged in equally by user and computer) rather, for example, than on the user's mental model. The computer is thought of as an enabling medium rather than a mere tool. Laurel's metaphor, expressed in terms of Aristotle's analysis of theatre, highlights the constructionist nature of interface design rather than the ontological properties emphasised in the 'mode' metaphor. Indeed, attributes at each level are constructed from those at the lower level. The agent is charged with the responsibility of building her own access to the digital environment. The insights gained by Laurel's approach seem mainly to have been applied to the design of interfaces that are meant to stay in their delivered form. A more recent, "dynamic", approach has been taken by computer manufacturers who recognise that many users want to configure their interface themselves (with scope ranging from the rather superficial choice of screen saver to more substantial matters of structure and mode of interaction). It seems to be more important to provide the user with a configurable interface than to provide a particularly elegant or efficient one: it is a consequence of the user's constructionist drive that the act of configuring one's own interface makes it preferable.

4.2 Open Source

The second logical step, after the construction of a personalised interface to the digital world, is the construction of digital entities that populate and interact in cyberspace. What should be the form of these entities? Along with use of the Internet and subsequently the web has come demand, from a surprisingly large number of users, for “open source” software. The average computer user interacts with an operating system by clicking on icons, dragging-and-dropping and so on. A user-friendly graphical interface (GUI) shields her not only from invoking commands directly (i.e. from typing the command name and whatever parameters it requires) but also, and more interestingly, from the underlying code that implements the operations. Consequently, even the experienced user has no way to access and modify the underlying source code, which executes operating system or applications commands. A system whose code is directly accessible to the user is said to be *open source*.

The high demand¹⁰ for open source code is a reflection of the number of users who prefer, where possible, the option of configuring their own software rather than making do with off-the-shelf packages. This provides further evidence for the strength of constructionism (quite apart from the other factors involved in supporting the open source movement, which include a feeling of ‘brotherhood’ made possible by the web and in opposition to being dictated to by a monopolistic software company), but there is also a new factor involved. The major “extraordinary success” of the 1990s was Linux, a free, open-source version of Unix, whose remarkable story provides evidence of what may be called *distributed constructionism*. To clarify the point, consider the difference between Richard Stallman’s and Linus Torvalds’ strategies.

¹⁰ A statistically insignificant presence in 1997, the popularity of Linux and the free/open source software movement has exploded in the last five years. In 2000, the International Data Corporation (www.idc.com) estimated that Linux was the fastest-growing server operating system, with 27% (up from 25% in 1999) of the server market, second only to Windows NT, which had 41% (up from 38 percent in 1999). Moreover, according to a new report from IDC (*Server Operating Environments Market Forecast and Analysis, 2000-2004*) commercial shipments of Linux will grow at a compounded annual growth rate of almost 17% from 1999 to 2004.

On the one hand, Richard Stallman's *Free Software Foundation* (begun October 1985, see Williams [2002]) released the code for components, as they were completed by Stallman himself, of his version GNU (GNU's Not Unix) of the Unix operating system (GNU/Linux). "The overall purpose is to give the users freedom by giving them free software they can use and to extend the boundaries of what you can do with entirely free software as far as possible." (Stallman, quoted in Moody [2002], 28). Stallman's GNU GPL (General Public License) perpetuates, efficiently, the freeness of open source software and any derivatives resulting from modifications by its recipients. "This enormous efficiency acted as one of the main engines in driving the free software projects on to their extraordinary successes during the 1990s" (Moody [2002], 28). Initially, circulation of the original components was by magnetic tape from Stallman or people affiliated to his project, when the web was not yet a common medium of communication. Controlled by Stallman, the enterprise still exhibited egopoietic values, most notably it was meant to promote a software version of the "freedom of speech" movement.

On the other hand, Linus Torvalds launched his project for the development of Linux by relying entirely on *distributed constructionism*, that is the unsuspected but evident interest, shared by a growing community, in coordinating efforts to achieve a global product whilst each realising only a local specific component of it. The project took full advantage of the web's *point-to-point* penetration. Human communities tend to be rigidly structured, so that direct communication between individuals is highly constrained. The media can be seen as partially facilitating that tendency, and mobile phones help to implement it to a restricted degree. But the web removes that constraint almost entirely amongst its "netizens", and provides a poietic-enabling environment through which the community of users and developers of Linux could interact and communicate easily and efficiently. Linux has clearly developed as an ecopoietic enterprise.

The difference between the two approaches has not passed unnoticed.¹¹ It has been well summarised by Eric Raymond in *The Cathedral and the Bazaar*: "Linux

¹¹ Moody [2002] seems to underestimate the "philosophical" contrasts between the two movements, on which see for the documents cited in the bibliography under the entries *Free Software Foundation Website* and *Open Source Software Website*.

overturned most of what I thought I knew. ... I believed that the most important software ...[like that of Stallman] ... needed to be built like cathedrals, carefully crafted by individual wizards or small bands of mages working in splendid isolation, with no beta to be released before its time. Linus Torvalds' style of development - release early and often, delegate everything you can, be open to the point of promiscuity - came as a surprise ... the Linux community seemed to resemble a great babbling bazaar of differing agendas and approaches". (Raymond [2001], 21).

The difference between Stallman's and Torvalds' strategies may appear to be partly attributable, historically, to different stages in the development of the Internet. Conceptually, however, it is really the result of two different constructionist ethics. Linux and other similar open source products are built and maintained as an expression of distributed constructionism on the web. They provide another dimension to Stallman's simple individual constructionism, and one supported amply by the web, which is thereby seen to provide a robust support for *collaboration without attrition*.

4.3 Digital Arts

The availability of web-based interfaces and software make possible the construction of forms of digital art previously unimaginable. Murray [2000] has identified three characteristic pleasures of digital environments in general:

- 1) *immersion*, the participatory immersive medium intensifies the age-old desire to live out fantasy. Rather than Coleridge's "willing suspension of disbelief", she proposes it to be viewed, more realistically, as supporting "the active *creation* of belief" (p. 110, emphasis added);
- 2) *agency*, that is 'the satisfying power to make meaningful action and see the results of our decisions and choices' (p. 126); and
- 3) *transformation*, that is the shape-shifting, morphing possible because of the digital representation of data and the ease with which it can be transformed.

For the purpose of analysing the future of digital narrative, Murray reflects: "These pleasures are in some ways continuous with the pleasures of traditional media and in some ways unique. Certainly the combination of pleasures, like the combination of properties of the digital medium itself, is completely novel". (Murray [2000], 181).

Murray's interest is in digital environments generally, not specifically in those supported by the web. The web is *public* in a way that other digital media are not. Nevertheless, if we add to Murray's three pleasures that of *interactivity*, we are led to investigate the wider field of digital art and the impact that constructionism has had on it.

Digital art has shared with CyberEthics its first half century of existence (Reffen Smith [1997]). Over this period, the topic has expanded with the pervasive influence of the digital medium and now includes graphic art, musical composition, poetry, architectural style and cinema as well as narrative fiction. Despite such variety, it seems that "digital art is novel in two ways, the first deriving from virtual reality techniques and the second deriving from the capacity of computers to support interactivity" (Lopes [2003]). Because the result of some digital art is difficult to distinguish from traditional art, emphasis is placed on the *process* rather than the *product*. (If a computer can solve crosswords faster than I can - albeit by the brute-force method of searching through a dictionary and trying all feasible combinations - then, one reasons, at least *the way* I do it cannot be mimicked by computer. Again, if a computer can produce Picasso-like pictures - albeit routinely by digitising a photo and then processing an abstraction of it - then, one reasons, at least Picasso's originality is inimitable). The same emphasis, on process rather than product, is made by Binkley (Binkley [1998]) who identifies the objects being manipulated, or *maculated*, by artists as being digital (data structures rather than paint or cardboard) with the result that the artwork produced lacks physical uniqueness and can in fact be copied electronically indefinitely. His view of process can be interpreted as acknowledging the importance of constructionism. Indeed, Binkley makes the point that, with the web, the objects of construction may bear little resemblance to those of earlier generations.

4.4 Homepages and the construction of the self

With interfaces, software and even new forms of art being constructed in cyberspace, the self is next in line. Web sites are certainly popular new objects of creation. The reason lies partly with the recent development of e-commercial models of marketing (if you want to buy a lounge suite, visit our web site and simulate how it would appear in your room), partly with human desire or need for communication (from government legislation

to photos of the grandchild's first birthday), partly with a new wave of constructionism concerning the self through personal homepages (Chandler [1998], see also Adamic and Adar [online]). Although graduates of computer science who once would have gone into programming jobs now go into web design, the sale of off the shelf software for constructing web pages is burgeoning.

4.5 Virtual communities

With the construction of the self, we have reached the starting point for the construction of virtual communities. What can we learn from socio-cyber-phenomena like web-based chat-rooms, interest groups, ICQ-like communities, newsgroups, online forum etc., which rely for their existence on point-to-point communication offered by the web? Until recently, it was common to argue, pessimistically, that the Web prompted people to withdraw from social engagement and become isolated, depressed and even alienated. According to a constructionist view, however, the Web actually provides a poietic-enhancing environment, which should facilitate, rather than hinder, the construction, the development and the reinforcement of self-identities, of links with local (real and/or virtual) communities and of social interactions. New data confirm this prediction. Virtual communities have become the most popular Web domain category, after search engines and portals (source: Nielsen//NetRatings, April 26, 2002, http://www.acnielsen.at/at/news/press/2002_04_26_110502/FULLTEXT.PDF). And a report published by the Pew Internet & American Life Project (<http://www.pewinternet.org/>), entitled *Online Communities: Networks that nurture long-distance relationships and local ties* (October 31, 2001), has shown that “the online world is a vibrant social universe where many Internet users enjoy serious and satisfying contact with online communities. These online groups are made up of those who share passions, beliefs, hobbies, or lifestyles. Tens of millions of Americans have joined communities after discovering them online. And many are using the Internet to join and participate in longstanding, traditional groups such as professional and trade associations. All in all, 84% of Internet users have at one time or another contacted an online group.” (http://www.pewinternet.org/reports/pdfs/PIP_Communities_Report.pdf). Virtual communities are a flourishing result of the free exercise of the constructionist drive. In

them, users reveal personal facts, “flame”, and switch personas by endlessly constructing, deconstructing and reconstructing alternative selves. They collaborate with and participate in a common social project. In general, they behave quite differently from the way they would behave in person. It is as if the normal metric of social distance were expanded by the web. The web empowers new categories of users with the possibility of constructing a new self and an e-polis. It makes constructionism an open option for anyone with access to an Internet connection.

4.6 Constructionism on the Web

What is the nature of constructionism as exhibited on the web? The previous examples show that the characteristic features of the web that seem particularly relevant to existing instances of constructionism are: interactivity, virtuality, agency, transformationality, process- (rather than product-) orientation, social publicity, and immediate point-to-point communication, which allows collaboration without attrition due to an apparent increase in social distance. Constructionism emerges as a most significant and intrinsic property of the web, more fundamental than any policy vacuum or pressing practical problems. The increased social distance means that the ethical consequences of constructionism on the web are particularly acute. Indeed, the apparent increase in social distance acts as a magnifier for ethical factors.

5. Homo Poieticus

Homo sapiens has primary needs, which relate to survival (like food, shelter, security and reproduction), and secondary needs (like hedonistic, intellectual, artistic and physical pursuits), which arise once primary needs are fulfilled. Constructionism seems to be amongst such secondary needs. It is the drive to build physical and conceptual objects and, more subtly, to exercise control and stewardship on them. It manifests itself in the care of existing, and creation of new, realities, being these material or conceptual. Thus, constructionism is ultimately best understood as a struggle against entropy. Existentially, it represents the strongest reaction against the destiny of death. In terms of a philosophical anthropology, constructionism is embodied by what we have termed elsewhere *homo poieticus* (Floridi [1999]). *Homo poieticus* is to be distinguished from

homo faber, user and “exploitator” of natural resources, from *homo oeconomicus*, producer, distributor, and consumer of wealth, and from *homo ludens* (Huizinga 1970), who embodies a leisurely playfulness devoid of the ethical care and responsibility characterising the constructionist attitude.¹² *Homo poieticus* concentrates not merely on the final result, but on the dynamic, on-going process through which the result is achieved. A punctured bicycle tyre may be mended entirely routinely (in primary fashion, for “survival” on a busy day) with little component of construction, or it may be mended in a more deliberate, considered fashion, perhaps with reflection on the process and what it is being achieved. In the case of the web, the ease with which digital constructs can be created and altered means that cyberspace is an ideal environment for *homo poieticus*. Many influential teachers of constructive disciplines emphasise in their teachings an approach to their art that we can now identify as constructionist, to distinguish it from the ludic, the routine or the mundane approach. Often these teachings draw from eastern philosophy and mysticism to make the point that the process, and the novice’s state of mind during it, are of fundamental importance. The end result will “take care of itself”, if the process is right.¹³

Given the importance we have attached to *homo poieticus*, it would be surprising if its nature had not been studied in other contexts. Two indicative examples are worth mentioning here, to enable the reader to place our position within a wider context.¹⁴

Piaget (Gruber and Vonèche [1995]) coined the term *constructivism* for an epistemic model in which children learn whilst interacting with their environment, by manipulating and building objects and developing coherent intellectual structures. Papert (Papert [1993]) extended Piaget’s work from genetic epistemology to the child’s construction of *microworlds* and called the result *constructionism*: “My perspective is

¹² Evers 2000 has associated the Open Source movement to *Homo Ludens*.

¹³ Particularly interesting examples of a constructionist attitude arise in most of the fine arts. To name just two, we refer to architecture (Alexander [1970], Liebeskind [online]) and cabinetmaking (Krenov [1976]).

¹⁴ We do not address here the critical issue of the connections between ethical constructionism and social constructivism. On the interactions between philosophy of technology and social constructivism see Brey [1997].

more interventionist. My goals are education, not just understanding. So, in my own thinking I have placed a greater emphasis on two dimensions implicit but not elaborated in Piaget's own work: an interest in intellectual structures that could develop as opposed to those that actually at present do develop in the child, and the design of learning environments that are resonant with them." Inspired by both, Murray (Murray [2000]) is interested, as we have seen above, in the possibilities for narrative fiction in Cyberspace. She uses Piaget's term "to indicate an aesthetic enjoyment in making things within a fictional world" (p. 294). Indeed, she claims that "constructivist pleasure is the highest form of narrative agency the MUD [Multi User Domain] medium allows" (p. 149). Whilst for Piaget and Papert the mental process of construction is autonomous and even subconscious, for Murray (and for us) it is typically explicit. More recently, constructivist methodologies have been applied to digital media. In Eisenstadt and Vincent [2000], for example, we read that: "Our approach to media rich learning experiences derives from constructivist models of education" (p. ix); the aim is "[...] empowering individuals to create their own content" (p. ix). In this case, the difference between the two approaches is that, for our constructionist perspective, the fundamental novelty brought about by computer-based or online learning has got little to do with long-distance courses, virtual classes and tele-presence, for it is rather to be identified in the vindication of the "maker's knowledge" tradition.¹⁵ ICT makes possible hands-on experiences, simulations, collaborations, and interactions with conceptual or information structures that can be built, manipulated, disassembled and so on, thus completely transforming the learning/teaching experience.

The process-oriented component of our concept of constructionism also has an interesting precedent in literary theory. *Genetic criticism (critique genétique)*¹⁶ was the name given in the early 1970s to an empirical approach to the literary act "d'expliquer par quels processus d'invention, d'écriture et de transformation un projet est devenu ce

¹⁵ The Maker's Knowledge Tradition goes back to Plato. It is the view that an epistemic agent knows, understands or otherwise epistemically controls better (or perhaps only) what the agent has made.

¹⁶ For a summary of genetic criticism and two case studies (Flaubert and Proust) see Schmid [1998]. There is an interesting tension produced by a rigid application of those ideas when text is interpreted as digital art; our notion of constructionism provides one resolution of it.

texte auquel l'institution conferera ou non le statut d'oeuvre litteraire.” (Gresillon [1994], 206). However, the concept differs from ours in subscribing firmly to written traces: “Genetic criticism has used the post-structuralist dissolution of the closed text to define its own notion of the fluid, dynamic manuscript text which, since it is not in any published form, is subject to constant revision. At the same time, genetic criticism has abandoned the vague post-structuralist conception of the text as an interactive process. The genetic approach reinstalls the text in its materiality. Its objects of inquiry are the material traces of writing.” (Schmid [1998], 12).

6. Conclusion: from CyberEthics to Information Ethics

For its first half century, CyberEthics, the ethics of ICT and in particular of the web, has been a situated action ethics. The point becomes clear if one reads Bynum’s overview (Bynum [2001], see also <http://plato.stanford.edu/entries/ethics-computer/>), which aims to survey the “historical milestones” of the subject, decade by decade. According to Bynum, “the best way to understand what the field is like is to examine some example sub-areas of current interest”. He considers the workplace, security, ownership and professional responsibility. Clearly, the approach to CyberEthics (Computer Ethics) has been predominantly pragmatic and action-oriented.¹⁷

In the absence of any foundational principle, the field is reduced to a collection of case-based analyses.

The battle cry for the 1990s has been James Moor’s quote: “A typical problem in computer ethics arises because there is a policy vacuum about how computer technology should be used” (Moor [1985]). In the tailwind of CyberEthics’s policy vacuum, much of the discussion has concentrated on the extent to which the web, or internet more generally, provides only a context of application for standard ethical issues *in silico* (the uniqueness problem of Computer Ethics). The conclusion has been that, at the very least, the web magnifies many ethical issues (security, privacy, ownership and so on).

¹⁷ For a conceptual analytic history of Information Ethics, complementing Bynum’s collection of milestones, see Floridi and Sanders [2002] and Tavani [2002].