

PREPRESS DRAFT of chapter in *The Handbook of Human Computation*, forthcoming from Springer New York. Expected publication winter 2014.

Human Computation in Electronic Literature

Scott Rettberg

The University of Bergen
email: scott.rettberg@lle.uib.no

Abstract: This chapter situates and considers several different facets of human computation in electronic literature and digital art. Electronic literature encompasses works in literary forms that are particular to the computer or the network context. Human computation is examined as an element of the development of collective narratives online, in which different roles are defined in architectures of participation. The form, structure, and common features of notable human-computation based artworks are identified. The human computation processes of collectively written and internet-harvested haiku generators are contrasted with each other to reveal their different models of situating the relationship between computational process and human authorship. Literary meta-critiques of human computation technologies such as Google's machine reading of Gmail and reCAPTCHA's use of human language recognition are discussed as electronic literature is positioned in a critical, if symbiotic, relationship to human computation.

1 Introduction

Louis von Ahn (2009) has described human computation as “a paradigm for utilizing human processing power to solve problems that computers cannot solve.” Quinn and Bederson (2011) further describe a consensus that what constitutes human computation are the problems that fit the general paradigm of computation, and as such

might be solveable by computers; and in which the human participation is directed by the computational system or process. A typical example of human computation would be an Amazon Mechanical Turk process using the incremental labor of internet workers to verify that images of red shoes for sale in an online store actually match the description of the product's color advertised on the site.

Most forms of electronic literature can be considered to have some elements of human computation: the majority of works in this field consist of texts authored by humans which are then subject to some sort of computational process or algorithmic manipulation. Electronic literature is a field of literary and artistic practice that, according to the Electronic Literature Organization, involves “works with important literary aspects that take advantage of the capabilities and contexts provided by the stand-alone or networked computer.” This encompasses a wide range of digital literary practices including hypertext fiction, kinetic poetry, chatbots, interactive fiction, interactive drama, generated poetry and narratives, narratives situated in networked communication technologies such as email, SMS, blogs, Twitter, and wikis, textual digital art installations, and many other practices. With electronic literature, human authors develop texts that involve computational processes—both texts that are themselves computer programs and texts that are the result of human interaction with algorithms—and human readers engage in reading practices that are technologically mediated.

Considering electronic literature from the standpoint of human computation is something of an inversion of the standard perspective. Scholars in this field more typically focus on how computers, networks, and computational processes can be useful in enabling humans to create new forms of literary expression, rather than beginning from the question of what roles humans play in a computational process. The challenges of creating a convincing and engaging narrative or producing a rich poetic use of language are still not generally solveable by computation alone. Even in the case of successful story or poetry generation, aspects of human writing are deeply integrated into the development of the system.

Hayles (2008) refers to the relationship between humans and computers evident in many works of electronic literature in terms of symbiotic loops: “Humans engineer computers and computers reengineer humans in systems bound together by recursive feedback and feedforward loops, with emergent complexities catalyzed by leaps between different media substrates and levels of complexity.” Likewise, the relationship between the system and the human participants / authors in works of electronic literature is often more complexly layered than a single iteration of enlisting humans to perform tasks the system cannot provide without human input. There are examples of works of electronic literature where human authorship is directed by computational processes. We encounter systems that are first developed—by humans—as literary platforms, which then computationally direct, arrange, or integrate contributions by other humans¹. The system may or may not be altered in response, in a recursive cycle that can continue.

After briefly discussing architectures of participation in collective narratives, I will focus herein on three types of human computation relevant to electronic literature:

- 1) Digital art projects involving human computation which offer some lessons for human-computation-driven electronic literature;
- 2) Poetry engines that use human contributions or human judgment to produce or refine combinatory or generate poetry;
- 3) Literary projects that are self-consciously engaged in a meta-level critique of the role that large-scale systems of human computation—for examples Google’s global-scale harvesting of search queries—play in reconstructing contemporary human culture and social practices.

¹ The *ePluribus Solver* project (Greene, Thompson, and Michelucci 2013) provides an example from the domain of collective journalism. Working with small fragments of a story in pictures using only a few characters or words, team members cast into descriptive and evaluative roles worked together to develop a collective narrative of the given situation.

2 Architectures of Participation: Frameworks for Collaboration

A literary project involving human computation should be understood to have an architecture of participation, a system that affords and constrains human participation. This architecture can be understood both as a platform in the sense of a computational system and a stage on which human interaction with the text, the system, and other authors and editors takes place.

Human computation in electronic literature is relatively uncharted territory. Paul Rohwer's "A Note on Human Computation Limits" (2010) considers two projects: *A Million Penguins*, a crowdsourced wiki novel produced by De Montfort University and Penguin Books in 2007, and two audio books produced by BBC Audiobooks America, that harvested Twitter responses to the first line of a story in "an iterative progression, singular integration model" to result in a collective fiction. The wiki novel project was an experiment in using the collaborative wiki platform—in which any user may edit any other user's text at any time (though those changes may be reverted)—to create collectively written novel. In their "A Million Penguins Research Report" (Mason and Thomas 2010) produced after the conclusion of the project, project organizers concluded that the result was ultimately less interesting as a novel than it was as a cultural text or performance. Penguin Digital Publisher Jeremy Ettinghausen reports "as the project evolved, I stopped thinking about it as literary experiment and starting thinking about more as a social experiment." Other critics and co-authors of the project recorded similar responses. The lightly controlled chaos of the wiki, it appears, served as a compelling arena for textual performance, but not for the development of a cohesive narrative.

Rohwer contrasts this project with one he considers successful, *Hearts, Keys, and Puppetry* by Neil Gaiman and the Twitterverse (2010). The story began with one tweet by Neil Gaiman, and readers then contributed Tweet-long continuations of the story. A single editor reviewed these tweets and selected the next line that would be included in the canonical version of the story, one line at a time. Rohwer argues that the "single real-time editor may be the natural

requirement to achieve a sufficiently coherent narrative.” While it is problematic to suggest that there is any “natural” requirement for coherent narrative—there are certainly many examples of multi-authored texts that did not have a single editor—it is clear that the two projects had different architectures of participation and control. The problem with narrative cohesion in *A Million Penguins* may have simply been that this architecture was not established as a system in which contributory and control roles were clearly defined and functional.

In a previous article focused on collective narratives (Rettberg 2011), I discussed a number of different online literary narrative projects that involved collaborative methods. These range from collaboration in small groups of authors, such as in the hypertext novel *The Unknown* (1999) to the attempt in the early 1980s by the Seattle writing group The Invisibles to use questionnaires and an early form of literary computer database to gather material for a novel, *Invisible Seattle* (1987), written by the whole city of Seattle, to projects such as Barbara Campbell’s *1001 Nights Cast* (2005)—a durational performance in which Campbell daily solicited individual texts from internet participants in response to a prompt which changed each day, and then performed a reading of one texts each night 1001 nights in a row. Surveying collective narrative projects, I identified three different types of participation:

Conscious participation: Contributors are fully conscious of explicit constraints, of the nature of the project, and of how their contribution to it might be utilized.

Contributory participation: Contributors may not be aware of how their contribution fits into the overall architecture of the project, or even of the nature of the project itself, but they do take conscious steps to make their contribution available to the project.

Unwitting participation: Texts utilized in the collective narrative are gathered by the text-machine itself, and contributors have no conscious involvement in the process of gathering the material.

Human computation-driven literary projects might involve any of these three different types of participation. People might be consciously participating as co-authors (for example by writing or editing a chapter of a wiki-based novel), may simply provide some text or information that will then be integrated by editors or by a computational system into a larger literary structure (for example respondents in the *Invisible Seattle* project who answered questions like “What is the best restaurant in Seattle to go for a break-up dinner?” and thus provided settings for the novel), or could be participating in a completely unwitting way (I will later discuss of Twitter haiku projects which harvest unwitting haiku from a general Twitter stream).

3 Digital Artworks Based on Human Computation

Electronic literature and digital art practices are deeply intertwined, so before moving to further specifically literary examples, it is useful to consider some notable examples of non-linguistic digital art that involve human computation. Aaron Kolbin’s “The Sheep Market” (2006) is a project that involved the production of 10,000 sheep by workers on Amazon’s Mechanical Turk. The workers were paid \$.02 for each sheep they produced. Kolbin developed a Processing-based drawing tool, which recorded the drawing of each sheep. Each worker was instructed to “draw a sheep facing left.” The results of the project included installations with prints of all 10,000 of the sheep, and animations, which reproduce the process of each sheep being drawn. Kolbin reports that the average wage paid to each worker was \$.69 per hour, and the average time spent on drawing each sheep 105 seconds.

One might reasonably ask what the point of such an experiment might be, or where we should locate the “art” in a project which is based very much on the idea of “amateur” production (albeit “professional” in the sense that each of the workers was paid). Certainly on some level there is an embedded critique of the labor dynamics of human computation. Paying someone \$.69 an hour for labor of any

sort is unconscionable by the standards of most developed nations². It calls into question other projects that use Mechanical Turk and similar platforms—is human computation simply a way of lowering labor costs to avoid paying human workers a reasonable minimum wage? And of course, the project also mirrors some more general global labor issues: Western consumers would not have access to such a plethora of affordable and wondrous consumer electronics without laborers in the East who are paid subsistence wages in poor working conditions to perform repetitive tasks. So on one level, the work can be understood as being about the political economy of contemporary consumer markets.

On the other hand, the process of human computation here also reveals tremendous creativity and diversity in a generalized class of human producers. Even in a simple rectangular black-and-white drawing environment, we encounter a diverse variety of approaches to producing a drawing of a barnyard animal. Like snowflakes, each of the 10,000 sheep in the market is in some way distinct from the others. The most fascinating aspect of watching the animations of the sheep drawings is seeing a human decision-making process unfold, as the workers draw, hesitate, make half-starts and scratch-outs. The drawings themselves are not nearly as affective as these ghostly presences, these invisible hands.

² In his contribution to this volume, “Labor Standards,” Alek Festiner begins to unpack some of the thorny conceptual and jurisdictional issues involved in utilizing a globally distributed casual labor pool for crowdsourced human-computation-based labor.




Fig. 1. Overview of “Seed Drawing 52” by Clement Valla. (Reproduced from the artist’s website)

Clement Valla’s “Seed Drawings” series (2011) likewise uses Mechanical Turk as an engine for a collective art practice. In this case, rather than being provided with a written instruction of what to draw, each online worker is provided with a “seed drawing”—a pattern—and instructed to reproduce it using a simple drawing tool. The results, the artist notes, are much like a game of “telephone.” The first drawing is placed on the center of a grid, and the drawings based on it appear adjacent to it as they are produced. As each worker produces a new drawing based on another copy, the variability also increases dramatically. So what, in “Seed Drawing 52,” for example, is seeded as a simple black-and-white line pattern might,

several generations later, evolve into an image of a face, or a coffee mug, or a letter, or a fish, or a star. As the original “message” is interpolated, its content changes significantly. One particularly interesting aspect of the drawings in “Seed Drawing 52” is that as the drawings are interpreted by different human actors, they generally appear to move from abstraction towards representation—at the center of the image we see abstract drawings but as we move to the outer parts of the grid, many more of the drawings are of recognizable objects or symbols. When charged with the pure task of mechanical reproduction, it seems the workers could not simply engage in automatic reproduction of the previous image, but were instead driven first towards interpretation. While a simple computer program could have replicated the seed drawing accurately in all 6560 squares, the human workers first reflected on *what they thought it was*, reproducing not the image but an idea of the object it signified, even if it may have originally signified nothing.




Fig. 2. Detail of “Seed Drawing 52” by Clement Valla. (Reproduced from the artist’s website)

Kolbin, Valla, and a number of artists have continued to explore this type of collective, human-computation-driven methodology in subsequent works. From the perspective of narrative generation, Kolbin and Chris Milk's recent 2012-13 project "This Exquisite Forest" is perhaps the most intriguing. In this case, each work begins with a seed animation: for example of a stick figure falling down at the beginning of "A Bad Day." A HTML5 web-browser-based tool then allows successive users to add new frames to the new animation. They might continue to build from the seed narrative, or they might build upon any of the resulting branches. The branching tree structure can be used in a number of different narrative or thematic ways. In some cases the trees are clearly based on continuing established narratives and taking a story to a new turn or diverted path, while in other examples the continuities are limited to those of visual style.

We can note common features in each of the three art projects discussed above that provide lessons for the production of successful literary works based on human computation:

- 1) In each case, the artists provide users with a simple tool and platform for developing their contributions;
- 2) Contributors are also provided with a clear and concise *constraint*;
- 3) While the constraint or instruction is explicit, the interaction of the user with the constraint is also the point at which *play* takes place in the system, as it involves a moment of interpretation and decision on the part of the contributor;
- 4) The essential element of what makes each work appreciable, as an aggregate, collective work of art is not the *accuracy* of the human response to instructions, but the *variability* of the human responses to the given constraints recognizable in the aggregate.

4 Online Haiku Generators Involving Human Computation

Many of the early experiments of net.art involved the aggregation of contributed texts by a number of different anonymous human actors. *The World's First Collaborative Sentence*, launched by Douglas Da-

vis in 1994 is one simple example of this. When reader / contributors open TWFCs in a web browser, they encounter a long unbroken stream of text, and a link to a web form which they can use to contribute to the work-in-forever-progress. The primary goal of the project appears to have been open performance on a global network—the instructions encouraged contributors to “WRITE, PERFORM, OR SING ANYTHING YOU WISH TO ADD IN WHATEVER LANGUAGE YOU LOVE TO THIS COLLABORATIVE WORK, JOINING HANDS AND MINDS WITH YOUR SISTERS AND BROTHERS OF WHATEVER RACE, REGION, OR BELIEF ANYWHERE IN THE WORLD...” Contributors were encouraged to add not only text but also “PHOTOGRAPHS, VIDEO, SOUND.” The only constraint was that the contribution could not include a period and therefore the sentence could theoretically go on forever³.

Like many net.art projects, TWFCs was largely about the early idealistic exuberance and utopianism with which many people took the Web as they first encountered it as a new medium for human expression. The possibilities of instantaneous publication with nearly global reach and the ability to share texts and collaborate with thousands of other people, without the intrusion of institutional gatekeepers, were still very new in 1994. The focus is largely on the novelty of the device and the medium itself. The project was successful insofar as its aim was to simply be a large-scale participatory text—more than 200,000 contributions were made to TWFCs between 1994-2000. But it would be difficult to assess its interest or merit as a literary work. When the goal of the project is unstructured participation, it is no surprise that the result was rambling and largely incoherent.

From the standpoint of human computation, more compelling examples of digital literature involve participatory structures that use human contributions in more specific ways, driven by constraints and


³ Davis’s work was live until the early 2000s when the scripts driving the project became non-functional in the context of the contemporary Web. In 2012, the Whitney Museum restored the digital work, releasing both a “restored” historical version and a fully functional live version which allows for new contributions.

processes intended to result in a coherent reading experience. These often involve the use of literary forms that are themselves constrained. Let us consider for example three projects that enlist human participation in the generation of online haiku.

Though the structure of the traditional Japanese haiku is more refined, in its English incarnation, haiku is generally understood to be a form of three lines in a 5 -7 -5 syllable structure. Haiku are often imagistic, and typically deal with two aspects of nature that when juxtaposed, can serve to startle the reader or bring about some sense of recognition. Given the comparative simplicity of the form in its English incarnation compared say to a Shakespearian sonnet, it is no surprise that it has been the subject of many experiments with combinatory, generative, or collective poetry. Haiku were in fact among the forms of some of the earliest experiments with poetry generation—in 1967 John Morris published “How to Write Poems with a Computer” describing his haiku generation program developed at Michigan State University. Morris both described his actualized program and conceptualized a better one that would balance an algorithmic process with elements of randomness, though, he confessed that he found the most affective poetry to be “...communication from a particular human being. And this is precisely what a computer is not.”

Nanette Wylde’s *haikU* (2001) is a project based on principles of user participation and on the use of a randomizing function to produce haiku that startle in the sense of producing *unintended* juxtapositions—no single author has determined which lines will appear together. The reading interface is a simple, spare web page. Every time a reader reloads the page, a new haiku is produced. Following a link to “Write haiku” individuals can submit their own haiku in three lines, each of which has its own button to post the line to bins of first, middle, and last lines. The poems delivered on each reload of the site are not the individual haiku as submitted by readers, but recombinations of these first, middle, and last lines of haiku pulled together in a variable way. Two reloads of the page produced for example “working round the world / the oven melting fire / brushed by a warm hand” and “under the rainbow / dew softly lays upon grass / hot sex in the night.” Reloading the page twenty times or so, it is

remarkable how many of the poems read as if they have been individually intended by a human intelligence. Most of the haiku, perhaps 80%, cohere quite well as poetry.



[Write haiku](#)

scenting crumbling stone
 reality augmented
 meaningless thinking

[about haiku](#)

© [Nanette Wylde](#) 2001-2009

Fig. 3. Example of a haiku. (Reproduced from the project site)

Wylde provides two opportunities for instructions to contributors. The first is on the brief “about haiku” page where she explains not just what the project is but what Haiku are: “Haiku traditionally reference a season and are generally observations of everyday life” and she attests that the “challenge of writing successful random haiku is that each line must be ‘open’ enough to create a connection with any two other random haiku lines. Successful random haiku develop an image in the reader's mind that gives cause for contemplation/reflection/awareness.” She reiterates these last two instructions on the “write haiku” page.

Write haiku

The challenge of writing successful random haiku, is that each line must be 'open' enough to create a connection with any two other random haiku lines.

Successful random haiku develop an image in the reader's mind that gives cause for contemplation/reflection/awareness.

First line = 5 syllables:

Second line = 7 syllables:

Third line = 5 syllables:

[Return to haiku](#) [Nanette Wylde](#)

Fig. 4. *Haiku* writing interface. (Reproduced from the project site)

In *haiku*, the combinatory form and structure of the project, in concert with the form and structure of the poetic form, and the fairly subtle instructions to contributors, lead to the production of a poetic database that works fairly well. While extremely simple in concept and execution, the combination of human-written lines and arbitrary structure results in new poetry neither completely determined by any human nor free of authorial intention.

Another online haiku generator project produced during the early 2000s, *HaikuTree.org*, attempted to bring human judgment to computer-generated haiku. Web readers would place generated haiku on a virtual tree. The haiku would be ranked by all these readers and would further “weather” over time. Only the most popular haiku would survive this process. In theory—though the project and its source code are no longer online—these selections would inform the process of refining the generator itself, to “help computers write better poetry.” It is unclear from the remaining project documentation whether by this the project developer meant that human judgment was directly informing and training the system via a machine learning approach or simply informing the human developer as he refined the system itself. In any case, poetry or story generators that are trained by human response to output are certainly conceivable as a branch of further research.

A number of more recent online haiku generator projects harvest human-written texts from the Internet, scan them for 17 syllable

count and appropriate word-breaks, break them into lines, and re-display them as haiku. One example of this is John Berger's @HaikuD2 Twitter account. In this case all of the text is human-produced but none of it is necessarily intended as haiku. It is only when Berger's bot provides line breaks and a #haiku tag that it becomes recognizable as such. The Twitter bot approach, at least in this iteration, may be more limited than Wylde's simpler system, which involves more intentionality on the part of the contributors. While some of the resulting haiku are clever or amusing in the way that they formalize language that is otherwise colloquial or banal, most of them simply read as tweets with line breaks, and not necessarily as particularly good poetry.

Based on a similar process to that of the Twitter haiku bots but generally producing more compelling results is *Times Haiku*. Developed by the software architecture staff of *The New York Times*, *Times Haiku* is driven by an algorithm that scans the text of articles published on the *Times* home page for potential haikus using a syllable count dictionary. The dictionary is regularly updated and modified by the *Times*' staff "with words like 'Rhianna' and 'terroir' to keep pace with the broad vocabulary of The Times" (Harris 2013). The algorithm discards haiku "if they are awkwardly constructed" (presumably meaning they don't break lines properly) and do not scan articles "covering sensitive topics" (presumably to avoid the production of deeply offensive haiku). Staff of *The Times* then read the haiku found by the algorithm. Human journalists who find a haiku "beautiful or funny or just a gem of a haiku" then select them for posting to a Tumblr blog. Selected haiku are posted by the system as an image file on the blog, and from there readers can share them on a number of social network sites. Each posting also includes a link to the original *Times* story. If the haiku produced by this process are not often imagistic or concerned with nature, they are often timely and amusing in their relation to contemporary culture. A couple of choice examples of haiku resulting from this process during June 2013 include: "There are horses who / can uplift, cause a chuckle, / spur a memory." (from June 11, 2013 story "Philotimo: A Horse Rescue Story") and "Young skin is spandex; / older is linen and needs / loving attention." (from June 4, 2013 story "'Counterclockwise' and 'Up' – In Pursuit of Longevity").

Consider the relationships between computer and human involved in the production of *Times Haiku*:

- 1) Human journalists write stories including lines which (presumably unwittingly) have the syllabic count of a haiku;
- 2) These are automatically fed into an algorithm which flags them as haiku;
- 3) The program's syllabic vocabulary is further modified by human actors;
- 4) Human curators then interact with a feed of texts that meet the basic formal requirements of haiku;
- 5) Selected haiku are then formatted by the system as image files and posted on a Tumblr blog;
- 6) Human readers then choose to share and propagate their favorite haiku.

Times Haiku provides a superb case of a recursive literary use of human computation. Without the computational system, the majority of the texts from *The Times* would never be recognized as haiku. Without the unwitting participation of human contributors, the texts would not exist at all. Without the conscious participation of human curators, the system would have a more limited vocabulary and would provide less aesthetically satisfying results.

5 Literary Meta-Critique of Human Computation

During recent years several e-lit authors have produced works that engage critically with human computation as an aspect of the contemporary network environment. In this case, the authors are not concerned as much with using human computation to develop collectively produced narratives or poetry, but instead with the systems of large international corporations such as Google and Facebook that regularly harvest and monetize information about their users and their behaviors on the network. Human computation is of course occurring on a large scale in these cases, as every time a user posts on Facebook or searches on Google, another contribution is made to a very large graph of extremely marketable information both about that specific user and about the broader contexts of human language

and society. A group of authors loosely centered on the Digital Language Arts program at Brown University have this taken as a particular concern and derived literary art from it.



Fig. 5. *American Psycho* recontextualized. (Photo reproduced from the project site)

Mimi Cabell and Jason Huff's *American Psycho* (2010) is a work that provides a context for considering how Google's different feedback mechanisms shape and control human experiences on the Internet. With this project, Cabell and Huff focused in particular on the Google Mail platform. They note "Google reads our mails, garners information from our personal messages, and uses that profiling strategy to select 'relevant' ads. It then displays those ads on the screen next to the very emails from which they were initially taken." In order to test the behaviors of this system, the authors chose to send the entire text of Brett Easton Ellis's novel *American Psycho* through Gmail one page at a time. They then collected the links that Google displayed, and printed a book, in which they left intact Ellis's chapter titles but eliminated the text of Ellis's novel, leaving only footnotes that recorded the links Google had provided for each

page of the novel. They report that some of the ads Google returned were directly relevant to the text from the novel—a scene in the novel involving the brutal stabbing of a dog and a man generated ads for knives and knife sharpeners—if at other times completely irrelevant to the context of the novel. Sections of the novel including racist language did not return any ads at all, indicating that Google’s technology has at least some censoring in place. Ads for Crest Whitestrips coupons were the most frequent single item to appear. The project might be described as a work of conceptual writing focused on revealing and foregrounding processes of human computation that we might take for granted in the course of everyday interactions on the network that simultaneously take advantage of us and make marginal but significant alterations to our communications environments.




Fig. 6. A page of Cabell’s and Huff’s *American Psycho* showing only references to advertising URLs generated by sending Brett Easton Ellis’s novel through Google Mail. (Photo reproduced from the project site)

Complex questions of who has—and who should have—access to shared literary heritage and linguistic data are at play in John Cayley and Daniel Howe’s *How It Is in Common Tongues* project (2012). They describe the overall project of *Common Tongues* as remediating “practices of and processes of reading” and critically addressing “the commodification of reading itself, and the proprietary enclosure of a growing portion of our linguistic cultural commons.” In particular the project addresses the fact that on the Internet many texts are now first read, processed, recomposed, and “multimediated” by computers in “pages that precede and predetermine any further or deeper ‘human’ reading.” The project, installed at the ELMCIP Remediating the Social exhibition at the Inspace Gallery in Edinburgh in November 2012, had a number of different digitally mediated text components that engage in different ways with the Google search engine, practices of reading, ownership of language, and Samuel Beckett’s work *How It Is*.

One aspect of Cayley and Howe’s installation notable for its engagement with copyright was a printed copy of Samuel Beckett’s text *How It Is*. While the text of the literary work printed in the book is identical on a word-for-word basis to Beckett’s text, every phrase in the book is footnoted with a URL. This URL corresponds to a non-Beckett use of the phrase found as a result of using a search engine. In his description of the project in the ELMCIP Knowledge Base, Cayley notes that all of the words in the book “are quoted from a portion of the commons of language that happens to have been indexed by a universally accessible engine.” Samuel Beckett’s estate, notorious for their enforcement of copyright, would doubtless have some issues with this citation practice. The work however makes the point that the text here is doubly enclosed: once in Beckett’s text by a copyright system that makes texts unavailable for reuse and adaptation until long after the authors are dead, and again as the texts that appear as search results by Google’s indexing system, which harvests texts written on the Internet by humans and machines and immediately begins making use of those texts everywhere it encounters them.

Samantha Gorman’s *Completely Automated* (2011) is an “exploration of how our written histories are forged through the interplay be-

tween human and machine editing.” The project engages critically with the human computation-based archival project reCAPTCHA—the system developed by Louis von Ahn which serves both as a spam blocker—by using language recognition to test whether a user is human—and as aid in the process of digital archiving of scanned texts—by using human responses to images of individual words in scanned archival texts to verify optical character recognition. Gorman produced a short film (2012) enacting a fictional scenario in which she can first be seen typing a text, “Pronouncement Against Domestick Production of Fraudlent Coinage as Decreed by Sovereign Law and Writ” by John Cartwright, into a page layout program, making modifications such as changing the name of the author, as she goes. She prints the modified text, outlines over the printed letters with painted ink, stains the paper with tea, giving it an aged appearance, before scanning the text into a university library’s archive system, and then planting it in a folder in the rare books room. The video concludes with other Internet users scrutinizing individual semi-observed words of the fraudulent text, as these fragments are approved one at a time.

Gorman explains the crux of her issue with the reCAPTCHA process on the project site: “Essentially, even a slight deviation from the original may escape the loop’s filters and be preserved digitally as a final authoritative text: our cultural heritage. Meanwhile, the original print is less conveniently accessible than the digital version and begins to lose authority within its physical library archive.” Gorman further suggests that, in privileging human language recognition, the reCAPTCHA system suggests that these processes are what “define us as human and . . . best distinguish human cognition from that of a machine.” So Gorman’s project raises conceptual issues with both the inherent uncertainty involved in integrating humans into computational processes—humans might not only make errors but conceivably could purposefully subvert the system—and with the effect human computation might have on the role and function of human cognition. Furthermore, in integrating steps of human cognition into processes that are controlled by machines, are we in effect subordinating human cognition, treating humans as superior sensory apparatuses, but lesser cognizers, than the machines they serve?

As the three projects discussed above reveal, the relationship between electronic literature and human computation is not simply procedural. While electronic literature authors may design architectures of participation to develop more effective collectively produced narratives, or new ways of harvesting poetry from streams of network discourse, they also have a role to play in critiquing the technological apparatus in which humans are increasingly embedded as actors, if not ghosts, in the machine.

6 Conclusion and Potential for Further Research

This chapter has considered human computation in a number of different aesthetic contexts: in the development of collective narratives, in massively crowdsourced visual and conceptual art, in haiku generators that automatically harvest and represent poetry from a Twitter stream or the news of the day. It has also considered how authors and artists are responding to a context in which their agency as creators or co-creators is resituated in relation to networked systems that are increasingly harvesting and interpreting human communications, reading and reformulating texts, and composing and determining narratives. The relationship of contemporary digital literary practice to human computation is neither entirely symbiotic nor essentially adversarial.

The field of electronic literature by nature experimental: practices from a number of different fields including writing, computation, visual arts, performance, communication, and design meet in this sphere. If there is a general commonality to the various practices and artifacts grouped under the rubric, it is that they all share an interest in exploring the relationships between literature and computation. It is important to emphasize that this a reciprocal set of concerns: we explore both the ways in which new possibilities for literature are afforded and constrained by computational processes and the networked environment and, in turn, the new possibilities for computation and the networked environment afforded by literary practice.

In the specific area of human computation and network-based collective writing projects, although there is a rich and growing body of experimental work in the area, a great deal of practical research remains to be done. Detailed analytic case studies are necessary to better understand how collective writing systems can best be harnessed to establish a level of aesthetic control and structure that would result in a sufficiently coherent reader experience while allowing for a degree of play, variability of response, and diversity of collective knowledge that could usefully enhance these sorts of projects and distinguish them from single-author literary endeavors. Our understanding of these practices would also be furthered by greater research collaboration between writers and artists working in electronic literature and digital art with computer scientists working in human computation, machine language learning, and other areas.

Given world enough and time, this chapter could have detailed many other extant experimental works of collective writing. It is a growing area of interest. Projects such as Judd Morrissey, Mark Jeffrey and the Goat Island Collective's 2007-2010 project *The Last Performance*, for instance, involved a collective narrative contributed to by more than 100 other writers, all responding to the same provided constraints. The short narrative and poetic texts they produced were then machine-interpreted, thematically cross-linked, and visualized in a number of different configurations. This deconstructed/reconstructed narrative architecture further served as a text and context for live performance⁴. Projects such as Brendan Howell's *Exquisite Code* bring algorithmic processes even more deeply into the writing process. In that project, a group of writers sit together in rooms writing for extended periods of time in response to prompts that they and system generate. The texts that they write are then periodically subject to "select/mangle" processes by the system. Each performance of this project so far has resulted in a book-length text which could be said to have been written both by the participating authors and by the machine itself, in what Howell refers to as a "c[ad]aver[n]ous exquisite_code life-work" (Howell 2008).

⁴ See Rettberg (2010) for further discussion of this work and strategies for reading *The Last Performance* as text and collective performance.

There are many questions we have only begun to address: of how to best make use of human computation strategies to develop compelling collectively written narratives, of how to integrate algorithmic procedures into writing processes in ways that produce aesthetically satisfying results, of how to productively integrate the artistic research strategies of electronic literature with the experimental methodologies of computer science, and indeed of how the function of literary writing in general changes in an environment in which networked systems are constantly harvesting and reframing texts of all kinds. We can only be certain that when confronted with technological opportunity, writers will continue to invent new literary forms and that contemporary literary works will continue to offer opportunities for reflection on the communication technologies, languages, and cultures of the era in which they are produced.

References

- Ahn, L. (2005). *Human Computation*. Doctoral Thesis. Pittsburgh: Carnegie Mellon University. <http://reports-archive.adm.cs.cmu.edu/anon/2005/CMU-CS-05-193.pdf>. Accessed 11 June 2013.
- Ahn, L. (2009). Human computation. In *DAC '09 Proceedings of the 46th Annual Design Automation Conference* (pp. 418-419). New York: ACM.
- Berger, J. @HaikuD2. Twitter stream. Accessed 11 June 2013.
- Cabell, M. and Huff, J. (2010). American Psycho. <http://www.mimicabell.com/gmail.html>. Accessed 17 June 2013.
- Campell, B. *1001 Nights Cast*. 2005. <http://1001.net.au/>. Accessed 18 June 2013.
- Cayley, J. and Howe, D. (2012). How It Is in Common Tongues. Documentation on the ELMCIP Electronic Literature Knowledge Base. <http://elmcip.net/node/5194>. Accessed 17 June 2013.

Davis, D. (1994). The World's First Collaborative Sentence. <http://whitney.org/Exhibitions/Artport/DouglasDavis>. Accessed 11 June 2013.

Felstiner, Alek. (2013). Labor Standards. *The Handbook of Human Computation*. Michelucci, P., ed. New York: Springer.

Gaiman, N. and the Twittiverse. (2010). *Hearts, Keys, and Puppetry*. Audiobook. North Kingstown: BBC Audiobooks America.

Gillespie, W., Rettberg, S., Stratton, D. and Marquardt, F. (1998) *The Unknown*. <http://unknownhypertext.com>. Accessed 18 June 2013.

Gorman, Samantha. (2011). Completely Automated. Project site. <http://samanthagorman.net/Completely-Automated>. Accessed 17 June 2013.

Gorman, Samantha (2012). Completely Automated. Video on Vimeo. <http://vimeo.com/33118528>. Accessed 17 June 2013.

Goodwin, D. (2000). HaikuTree.org. Documentation of project at Rhizome.org. <http://rhizome.org/artbase/artwork/2178/>. Accessed 11 June 2013.

Greene K., Thomson, D., Michelucci, P. (2013). Explorations in Massively Collaborative Problem Solving. *The Handbook of Human Computation*. Michelucci, P., ed. New York: Springer.

Hayles, N.K. (2008). *Electronic Literature: New Horizons for the Literary*. Notre Dame: University of Notre Dame.

Harris, J. (2013). About Times Haiku. <http://haiku.nytimes.com/about>. Accessed 12 June 2013.

Howell, Brendon, et al. (2008). Exquisite Code. <http://exquisite-code.com/>. Accessed 21 June 2013.

Invisible Seattle: The Novel of Seattle, by Seattle. (1987). Seattle: Function Industries Press.

Kolbin, A. (2006). The Sheep Market.
<http://www.aaronkoblin.com/work/thesheepmarket/index.html>. Accessed 11 June 2013.

Kolbin, A., Milk, C. (2012). This Exquisite Forest.
<http://www.exquisiteforest.com>. Accessed 11 June 2013.

Quinn, A.J., Bederson B.B. (2011). Human computation: a survey and taxonomy of a growing field. In *CHI '11: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1403-1412). New York: ACM.

Mason, B. and Thomas, S. (2008) “A Million Penguins Research Report.”
<http://www.ioct.dmu.ac.uk/documents/amillionpenguinsreport.pdf>. Accessed 17 June 2013.

Morris, J. (1967): “How to Write Poems with a Computer.” *Michigan Quarterly Review* 6.1 (pp. 17–20).

Morrissey, J., et al. (2007). *The Last Performance [dot org]*. Accessed 21 June 2013.

Rettberg, S. (2010). Performative Reading: Attending *The Last Performance [dot org]*. *Dichtung Digital* 40. <http://dichtung-digital.mewi.unibas.ch/2010/rettberg/rettberg.htm>. Accessed 21 June 2013.

Rettberg, S. (2011). “All Together Now: Hypertext, Collective Narratives, and Online Collective Knowledge Communities.” In *New Narratives: Stories and Storytelling in the Digital Age*. Page, R. and Thomas, B. eds. Lincoln: University of Nebraska Press.

Rohwer, P. (2010). *HCOMP '10: Proceedings of the ACM SIGKDD Workshop on Human Computation* (pp. 38-40). New York: ACM.

Times Haiku Generator. (2013). <http://haiku.nytimes.com/>. Accessed 12 June 2013.

Valla, Clement. (2011). Seed Drawings.
<http://clementvalla.com/category/work/seed-drawings/>. Accessed 11 June 2013.

Wylde, Nanette. (2001). haikU.
<http://preneo.org/nwylde/haikU/index.shtml>. Accessed 11 June 2013.