Computational Contributions to the Humanities*

John Nerbonne
Rijksuniversiteit Groningen, The Netherlands

Abstract

At the University of Groningen we have emphasized a simple view of humanities computing as computing in service of the humanities. This means that we seek to answer scholarly questions in linguistics, history, and art history by using the computer, exploiting especially its ability to process large amounts of data and the transparency of its processing. We have shied away from questions of digital culture, avoided overemphasis on pedagogical applications of computers, and eschewed visions of scientific revolution—including, in particular, the revolutionary idea that humanities computing is a discipline, preferring to think of it instead as a federation of disciplines, whose practitioners find it opportune to collaborate for reasons of some common problems. We have discovered that our ability to deal with large amounts of data marks the distinctive contributions we can make to humanities scholarship.

Correspondence:
John Nerbonne,
Rijksuniversiteit Groningen,
9700 AS Groningen,
The Netherlands.
E-mail:
nerbonne@let.rug.nl

* This paper under the title 'Date Deluge', was delivered as the keynote plenary address on 11 June 2004 at the joint conference of The Association for Literary and Linguistic Computing and The Association for Computers and the Humanities at the University of Gothenburg. I appreciate the invitation to speak and several astute remarks on the talk, in particular from Patrick Juola, Geoffrey Rockwell, and John Unsworth. Thanks also to Koen De Smedt, Willard McCarthy and Lisa Lena Opas Hänninen for generous comments on an earlier draft.

1 Introduction

In Groningen we have emphasized a straightforward view of humanities computing: we seek to further humanities scholarship using the computer. In fact we have largely emphasized linguistics, but we also work in history and art history, and we have recently initiated a research line in communication. We have consistently aimed to further scholarship in these fields using computational techniques, and we can point to a number of results, i.e. contributions we have made to these fields. This paper suggests that humanities computing should focus on contributing to humanities scholarship, and that other goals, e.g. data annotation, applications, developing a new academic discipline, reflecting on common methods, or improving pedagogy—interesting and worthwhile as they may be—are at best derivative, and at worst (when they do not follow from the primary goal of furthering humanities research) distracting, or self-defeating. We also suggest that projects requiring the processing of large amounts of data are those where humanities computing is poised to make the greatest contributions.

doi: 10.1093/llc/fqh041

1.1 Motivation

It would be great to be able to point to some consensus among our non-computational colleagues about the significance of our contributions, including our (likely) future contributions. These will ultimately be used as a gauge of our value, and will in turn determine the chances we will have at future success. We could hope for some consensus about our worth if we could show progress in linguistics, history, and literature enabled by the computer. We submit that no such consensus exists today.

Computational techniques have been poised to contribute to humanities scholarship for some time, and yet, large numbers of significant contributions are not widely known. This is already unfortunate. Strategic discussions about humanities computing often revolve around preparatory activity on the one hand and data archiving and annotation, and millenary visions on the other and focus on potential for methodological innovation and refocus of scholarly aims. These activities are worthwhile, but only to the degree to which they contribute to furthering humanities scholarship, or are motivated by reflection on it.

We trace below how the pursuit of progress in the humanities leads naturally to methodological reflection and a redirection of scholarly aims, arguing that these follow naturally in the course of scholarly reflection. Focusing on redirection without building on results puts the cart before the horse. This paper is a plea for a focus on results in humanities computing.

The following section illustrates the sort of concrete contributions we suggest need to be emphasized. We note that common to all of these contributions is the use of the computer to analyze large amounts of humanities data, e.g. data about trade history, data about the sorts of constructions in use in a language, and data about regional pronunciations and regional vocabulary.

This is followed by some remarks on the nature of humanities computing, a plea to emphasize results rather than philosophy in seeking to advance the field, some notes on other views, and finally some conclusions.

2 Example Results in Humanities Computing

This section provides some concrete examples of results we have achieved in Groningen, all of which have involved well-established research questions in humanities and the processing of large amounts of humanities data, i.e. much more than would have been feasible to process reliably by hand.

Welling (1998) analyzed the import register of ships arriving at the port of Amsterdam in the period 1771–1871. He needed to digitize a substantial and difficult data source: standardize the names of ships, cargo types, and ports of origin; organize the information into a database; and finally, compare the results of his analyses with those of other historians. Even though Baltic trade is known in Dutch economic history as the *moedernegotie* 'mother trade', Welling could show that American trade had eclipsed the Baltic in value as early as 1780. Although not a focus of his analysis, Welling could also show that the Anglo–Dutch wars effectively eliminated the Dutch from competition with the English in shipping, and thus created a need for shipping which American vessels quickly filled, occupying second-place in shipping worldwide as of 1820.

Koster (2001) applies image processing techniques, in particular, so-called REGISTRATION techniques to align historical city maps and coordinate the information in them. His work interests students of architectural history who wish to obtain more comprehensive views of earlier cities, in order *inter alia* to understand how architecture and art was viewed in context. In this case the research has not (yet) involved large numbers of maps, but the procedures to align different maps make use of a great deal of information in the maps, and cannot be done effectively by hand.

Malouf and Van Noord (2004) apply training algorithms from machine learning in order to train preferences in a linguistically inspired syntactic parser to provide analysis trees of sentences in Dutch newspapers. Van Noord (2004) applies the software developed (as well as some clever heuristics designed to recognize gaps in coverage) to several million sentences (≈50 mil. words) in order to check the reliability of the syntactic analysis.[1] This—together with some clever tools implementing the heuristics (van Noord, 2004)—has resulted in several corrections of the best linguistic descriptions available for Dutch, e.g. corrections concerning alternative genders, aspectual auxiliary choices, case analyses, and patterns of complementation. This information would never have been collected without systematic computational examination.

Gaustad and Bouma (2002) apply text classification algorithms to incoming emails in Dutch at the service desk of a company specializing in providing Internet access. In order to improve (subject) classification, they first 'stemmed' the texts using software they developed. The stemming process reduced different forms of the same word to the same stem, e.g. *robs, robbed* → *rob*, but without guaranteeing that the resulting stem is the corresponding dictionary entry (so that *ride, riding* → *rid* also occurs). It can be thought of as a simple but not entirely correct version of lemmatization. The results indicate that stemming does not consistently improve classification, which is interesting for such practical applications of text classification, but it is also interesting linguistically. There are more instances of each word (stem) than there are instances of each of its separate inflected forms, so that the frequency of putative indicators rises, which should improve the reliability with which topics can be diagnosed. But no such effect could be demonstrated.

Kleiweg *et al.* (2004) look at the problem of projecting the information resulting from clustering to a geographic map (see Fig. 1).

1 Largely because of van Noord's computationally demanding work Humanities Computing is the second largest user at our university's High Performance Computing Center.

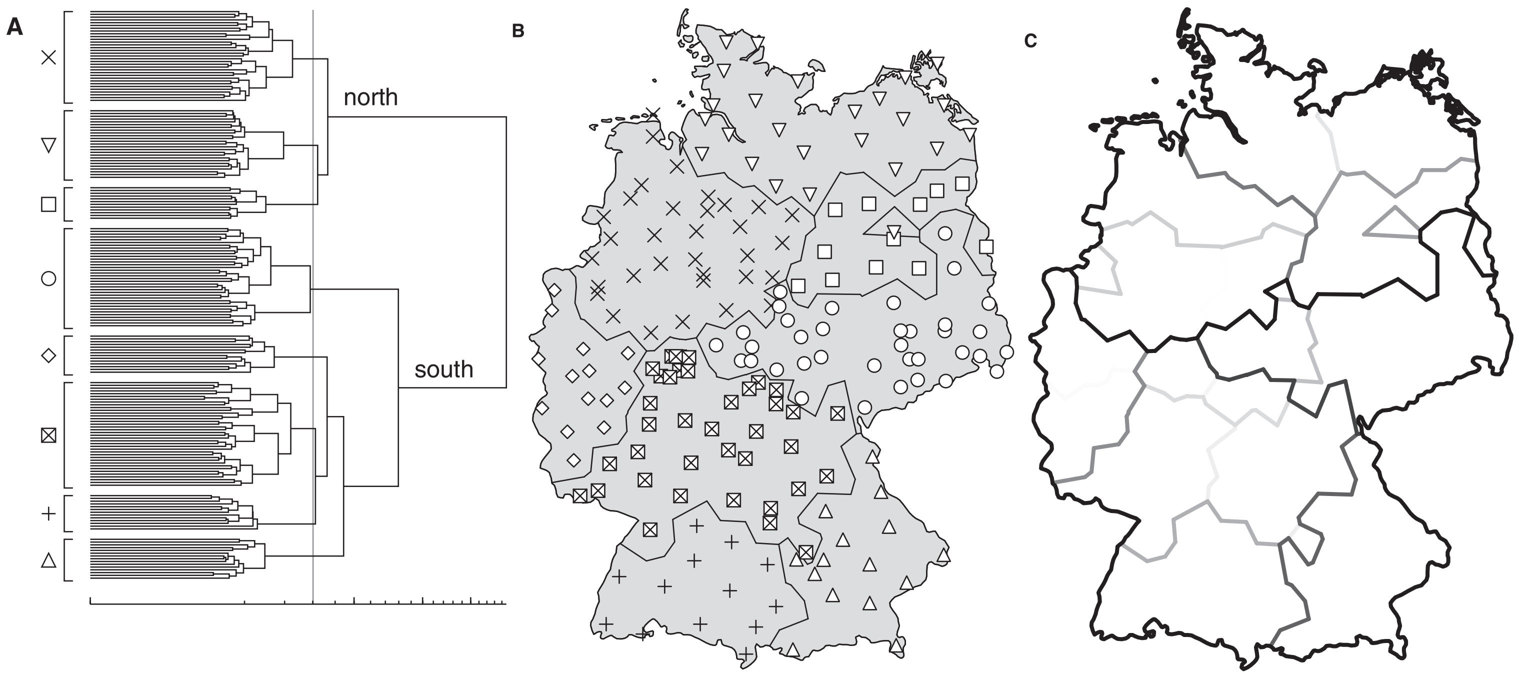

Fig. 1 Normally one simply projects a fixed number of groups from a clustering dendrogram (A) onto a number of discrete areas in a map (B), hiding a good deal of structure. A composite cluster map (C) is obtained by drawing every border implicit at some level of classification in the dendrogram, allowing the darkness of the borders to reflect the frequency with which the border is drawn. See Kleiweg *et al.* (2004) for further advantages

They produce not only maps, which dialectologists (and others) can inspect in order to understand the classifications inherent in their data (and their techniques), but also freely available software that performs the same analysis on other data to which clustering may be applied. This is one product of a new research line in Groningen, aimed at understanding graphical communication.

2.1 Proceeding from linguistics

Groningen Humanities Computing is focused on computational linguistics, which is a well-established interdisplinary field with a strong tradition of serious computational work on language. Computational Linguistics has an international professional society with 1,500 members who meet roughly twice a year in conferences with 300–750 participants (www.aclweb.org). In general it has been enormously beneficial to begin from this point of strength in Groningen, but it has also meant overcoming the prejudice that linguistics and computational linguistics are so *sui generis* that other humanities benefit little from their examples. Some colleagues view computational linguistics in particular, more as an applied (engineering) field, and see little linguistic advance in it, and others lament the fact that linguistics and philology have grown apart in the past half century. We do not agree with these complaints, but it is nonetheless true that work in computational linguistics requires—in addition to linguistic expertise—a high degree of technical skill. And it is also true that while a great deal of linguistics sees itself as allied primarily with cognitive psychology, there is nonetheless a great deal of interest in topics involving culture, especially social culture and anthropology.

We examine one project in more depth in order to substantiate our claims that some of the more ambitious goals others set for humanities also come within range of a program emphasizing research results.

2.2 Dialectology

Dialectology is an area in which we were able to proceed from (linguistic) research strength, but where we felt research might engage our other colleagues in humanities more seriously. Accordingly, we set out on an initially modest program applying some well-known techniques from computational linguistics to seek answers to important research issues in dialectology. The question we focused on initially was the nature of dialect areas. These have been a source of analytical puzzlement. While researchers tend to agree when they divide a language area into various dialect areas, still a precise formulation inevitably seems to stumble on the fact that individual distinctions often do not overlap perfectly, e.g. whether the word *light* is pronounced with a diphthong as in the Northern US [lait] or a monophthong as in the South [lat], or similarly whether the initial sound in *afternoon* is pronounced as a diphthong ([æᵊftənʉn]), South) or as a monophthong ([æftərnun], North). The imperfect overlap of

	Operation	Cost
æəf t ən ʉn		
æf t ən ʉn	delete ə	d(ə,[ ])=0.3
æf t ər n ʉn	insert r	d([ ],r)=0.2
æf t ər n u n	replace [ʉ] with [u]	d([ʉ],[u])=0.1
	Total	0.6

Figure 2. Levenshtein distance between two sequences is the least costly sum of costs needed to transform one string into another. The transformations shown here are associated with costs derived from phoneticians' work on the distance between individual phonetic sounds. The pronunciations are from the *Linguistic Atlas of the Middle and South Atlantic States* (Kretzschmar *et al.*, 1994)

dialect features, even closely related ones, is noted as standard in all of the handbooks (Chambers and Trudgill, 1998, [1980], p. 38).

There are several ways in which linguists have tried to measure the distance between two basic sounds, most of which are based on the description of sounds via a small (≤25) number of features (see Heeringa (2004) for details). There is also a standard technique for the computational comparison of sequences, namely, Levenshtein distance, also known as edit distance or string distance, and we set out to combine these techniques. The basic idea behind Levenshtein distance is to imagine that one is rewriting one string into another. The rewriting is effected by basic operations, each of which is associated with a cost, as illustrated in Fig. 2.

The operations used were: (i) the deletion of a single sound, (ii) the insertion of a single sound; and (iii) the substitution of one sound for another. Other operations are possible. The operation costs used in the procedures were those assigned by phoneticians (again, we refer to Heeringa (2004) for details). They consist of the measure of the distance between the sounds (in the case of substitution), and the measure of the distance between a given sound and silence (in the case of insertions and deletions)

If we apply a Levenshtein procedure to about 100 words from several hundred field work sites, the result may be shown to verify the idea of dialect areas as used in traditional dialectology (Nerbonne *et al.*, 1999). These may be reconstructed via clustering techniques, but also via the statistically more stable multi-dimensional scaling.

2.2.1 *Analysis*

Naturally, this (apparent) success led to a number of questions. We were first confronted with the statisticians' question (on 'data snooping'): had we demonstrated anything, or had we merely been lucky (or assiduous)? The procedure was sketched roughly above, but a great number of details can be adjusted, e.g. the treatment of diphthongs as one or two segments, the treatment of length, the particular segment distances used, the attention paid to frequency, the status of words as opposed to individual sounds, whether a ceiling should be placed on segment distances, etc. Similar questions arise in any number of computational fields due to the flexibility and power of computational modeling. How can we know that we are doing more than fitting a model to a random data set? We are pushed to this reflective step given the (perhaps uncertain) results obtained.

To answer these questions, we performed some meta-analyses. First, we could show that the individual Dutch words, as single probes of dialect distance, showed an inter-item correlation of $r = 0.19$ (while for Norwegian words $r = 0.10$). From this we can derive Cronbach's α measure of consistency, where $\alpha = 0.97$ for 125 Dutch words (and $\alpha = 0.86$ for 58 Norwegian words). We clearly have enough material.[2] To measure validity, we first examined overlap with consensus expert opinion (Heeringa *et al.*, 2002), where we obtained quite high measures of overlap between the expert classification and the classification obtained via Levenshtein distance in conjunction with clustering. However, we wished to avoid simply reaffirming scholarly tradition in our validation work, so we sought alternatives to consensus expert opinion as measures against which to judge the work. Again, the standard scholarly habit of questioning results pushes us towards further reflection.

Heeringa and Gooskens (2003) suggest a solution to the problem, namely that one obtain assessments of dialect distance from lay dialect speakers via experiment, and that one should validate computational methods *inter alia* through the degree to which they assign distances which correlate with the judgements of lay dialect speakers. This step is surprising from the point of view of standard dialectology, which often distinguishes between 'objective foundations' for dialect classifications, such as the pronunciation of the vowel in *light* and 'subjective foundations,' i.e. the judgements of lay dialect speakers. Heeringa and Gooskens effectively suggested using the subjective results in order to validate the objective ones! On deeper reflection, however, it seems clear that there can be innumerable 'objectively determined' dialect maps, as many as there are ways of analyzing language differences. Heeringa and Gooskens have thus contributed not only to the computational aspects of dialectology, but also to the foundations of the field.

2.3 Carrying on

Given a consistent and valid technique for measuring the distance between dialect variants, there are any number of interesting questions that can be asked. Heeringa *et al.* (2000) investigate the dialects spanning the Dutch–German border in Bentheim, noting substantial convergence toward the standard languages—the varieties on Dutch soil are converging toward standard Dutch and the varieties in Germany toward German. Heeringa and Nerbonne (2000) compare Dutch language varieties in their development over a period of 120 years, Gooskens and Heeringa (2004) examine distances between related languages and pose questions on mutual intelligibility, and

2 $\alpha = 1.0$ indicates perfect reliability. 0.6 (sometimes 0.7) is required for scientific purposes in psychology, and 0.9 for purposes with significant social consequences, such as educational testing and school admissions decisions.

Nerbonne and Kleiweg (2003) apply a simpler measure to lexical distance among American varieties. We are continuing this line of research vigorously.

We are also now in a position to ask deeper questions. Heeringa and Nerbonne (2002) suggest a solution to an old conundrum in dialectology. Suppose a traveler began at one end of a language area and walked across it, stopping regularly at towns and villages. The traveler would notice gradual differences at each stop along the way, but (on many lines through language areas) no abrupt changes. Nonetheless, the traveler might have walked through different dialect areas. Heeringa and Nerbonne compare incremental and cumulative changes from site to site along a path, and show that the traveler who only notices incremental changes is overwhelmed with unsystematic variation. The measurements bear this out. Unsystematic variation dominates in measurements over short distances.

We are likewise in a better position to ask explanatory questions, since we can now approach these with numerical techniques. These might be questions about the relative contributions of pronunciation, lexical choice, and grammar to dialect distance, or perhaps the extralinguistic determinants of dialect distance. Chambers and Trudgill (1998, [1980], Ch. 11) suggest that social contact keeps language varieties alike, and they note several reasons to suspect attractive forces at work here. Since the chance of social contact might be expected to drop as the square of distance, much as the force due to gravity does, the idea also bears the name 'gravity theory.' As Fig. 3 indicates, however, the dominant tendency in the data does not confirm this view.

This examination of the relation between geography and dialect distance employed distance 'as the crow flies,' a notion which might reflect the chance of social contact on an ideal plane, but which might be importantly refined. Van Gemert (2002) takes up the question of more refined geographical models in which travel time is estimated, and Gooskens (2004) and Gooskens (N.d.) show a major improvement in analysis when travel time rather than direct distance is used as the predictive variable in Norwegian data. Norway's mountains and fjords distort 'as-the-crow-flies' estimates of travel time a great deal.

2.4 Radical innovation?

As the last section documented, successful efforts in humanities computing can lead to radical reflection on foundational questions within a (sub)discipline and also to the generation of research questions that would have been impossible earlier. Given the urgency of obtaining good research results in humanities computing, and the likely ties to foundational and novel questions, we suggest that a focus on strong research results is a sufficient agenda. If we conscientiously apply computational techniques to humanities questions, the innovations, and also the reflection on foundations, will arise on their own.

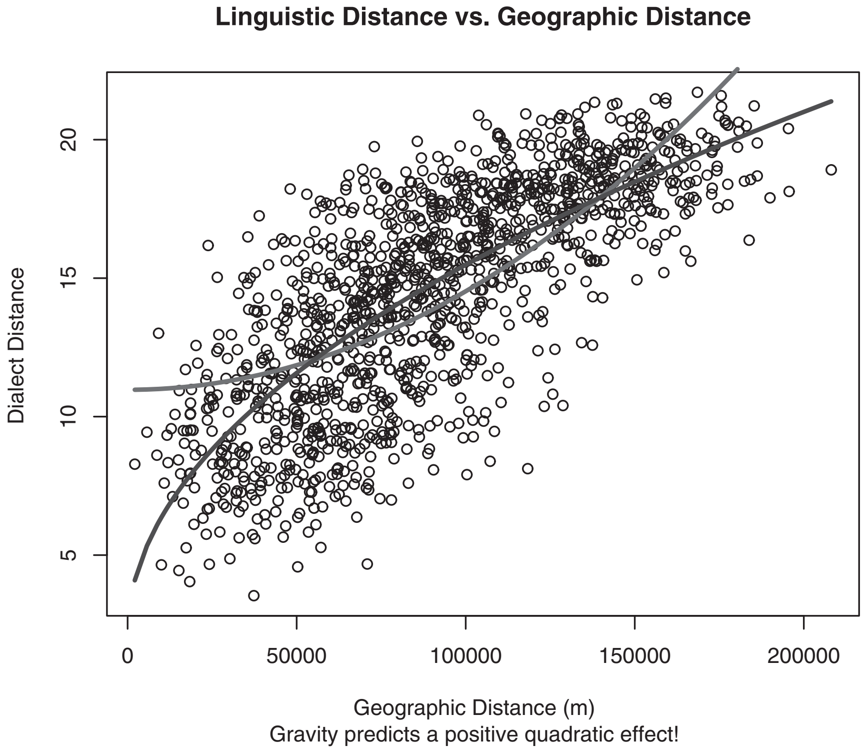

Fig. 3 'Gravity' models of social influence predict that linguistic distance should increase with the square of geographic distance (the rapidly rising higher line). While the correlation is certainly positive, as 'gravity' predicts, linguistic distance seems rather to grow with the square root or logarithm of geographic distance rather than its square

As we examine research possibilities in Groningen, and as we follow our own developments, we do not seek focus either on foundational issues or on the question of whether radical innovation is likely to ensue. We find that these arise naturally. Our key questions are in Fig. 4.

3 What is Humanities Computing?

In Groningen we view humanities computing as a federation of disciplines and subdisciplines, one in which practitioners face enough common problems to warrant collaboration. Our staff members are linguists, historians, art historians, and communication experts, all attempting to use the computer for their own purposes. The disciplines and subdisciplines united in Groningen's department are computational linguistics, historical information science (Boonstra *et al.*, 2004), architectural history, and communications. In the case of architectural history, there is not even a recognized computationally oriented subdiscipline, but there is interesting work, which is crucial.

Naturally, humanities computing could develop into a discipline of its own. Hopcroft (1987) reminds us in his Turing award lecture that

1. What substantial humanities question is posed? Who cares about the answer?
2. Is the computer essential? Is lots of data involved?
3. Have we really answered, or have we just tinkered for a long time?
4. If the project is successful, what will we be able to do?

Fig. 4 Crucial questions for projects in humanities computing

computer science, too, went through a period where its subject matter was taken from various other disciplines, including electrical engineering, mathematics, and linguistics, and in fact, he praises the early advocates of computer science for their willingness to promote the field long before they could define it. But things could just as well develop differently in the humanities.

We recognize a scientific or scholarly discipline by: its subject matter, normally a particular perspective on objects or events in the world; a body of theory about its subject matter; some analytical techniques with which new questions may be addressed; and, where appropriate, some practical applications. With this view, solid-state physics, Dutch history, and phonetics are reasonably regarded as disciplines (or subdisciplines), but not, crucially humanities computing, which has neither coherent common subject matter nor common theory. The last remark may sound harsh, but De Smedt (2002), summing up the results of a European survey of educational programs in Humanities Computing, noted the substantial variation (§ 4) in these programs, which reflects the lack of consensus on the nature of humanities computing among curriculum designers.

It makes sense to keep this in mind because we may identify common theory and technique if we are sensitive to the question. But a focus on seeking these commonalities may not further humanities computing (*pace* Mccarty (1998), see below). It would distract us from more important work in advancing the humanities through our research. In this we agree with De Smedt (2002, pp. 96–97), who calls on curriculum designers to keep research advances in mind. The urgent issue is to prove the value of humanities computing through results in humanities fields.

4 Emphasizing Results

Our simple conception of the proper goal of humanities computing has several advantages. First, the value of the humanities can be assumed—we do not need to remap the division of intellectual labor to make room for any new subject matter, and the value of the humanities is well accepted. Naturally, one can insist on scepticism with respect to these points, just as on most, but that would be a philosophical exercise.

Second, we approach colleagues in questions of resources and, ultimately, continued institutional existence in the most dignified way we know, by asking that we be judged on the basis of contributions to our common fields. Belying claims of Luddite attitudes of humanities scholars, our experience is that colleagues are open to new techniques if they first are shown reason to trust them—interesting results.

A third reason is that we are ultimately judged on our results. Humanities computing is much too mature as an academic activity to plea that it needs time to develop properly. In fact, humanities computing is rapidly approaching middle age. The journal *Computers and the Humanities* is scheduled to publish volume 39 in 2005 (as this goes to press). It is completely reasonable for a new field to ask for time in order to prove itself, but humanities computing has already had a comparatively long grace period. It is now time to produce.

A fourth reason to advocate emphasizing sober advances in science and scholarship concerns the nature of humanities computing. Any research question we seek to answer using computational methods is going to involve considerable innovation in methodology, argumentation and in most cases, perspective. In other words even if we aim to progress soberly, we will inevitably be innovative in all sorts of ways. These factors in and of themselves impede acceptance. Aiming at radically novel forms of scholarship or recognition for these will only exacerbate difficulties.

It has not escaped us that the view we emphasize is conservative in its assumptions about science and scholarship. Rather than urge a revolutionary view of the role of computing and the humanities, we advocate that we demonstrate the value of computing approaches in areas of undisputed interest. This is challenging enough, and there are enough areas of innovation in methodology and interpretation to satisfy all but the most impatient innovators.

Finally, a personal note. Like many similar departments, alfa-informatica (www.rug.nl/let/informatiekunde) in Groningen has suffered storms, but we believe that our focus on results in humanities disciplines has helped us weather them. We should note that there were once five similar departments in the nine Dutch universities (that have humanities), and that we are unfortunately the only surviving department. We believe that our adherence to this simple conception of the proper goal of humanities computing may have helped.

5 Other Views

The advantage of simple views is their relative clarity, which is, however, coupled with the dangers of vacuity and reductionism. Is it then simply tautologous to regard humanities computing as the application of computing to research questions in the humanities? Are other views even possible? In examining alternatives we proceed from the assumption that all humanities computing colleagues approach their work with special computational expertise, and that this is what sets us apart.

Humanities computing is not data annotation, not even very sophisticated text annotation such as SGML or TEI. We have undertaken work in this direction in the past in Groningen (Bouma and Kloosterman, 2002; van der Beek *et al.*, 2002), and we will undoubtedly wish to continue it in the future, but we always ask what gain in scholarly knowledge may be expected from the annotation effort. Annotation is not a goal in itself, not even when attention is focused on the annotation scheme, or on meta-systems for designing such schemes. All of this is preparatory.

Humanities computing is not the study of digital culture—even if we may wish to exploit our affinity with digital culture in engaging our colleagues in humanities (Nerbonne, 1998). The particular goal of understanding culture is naturally one which the humanities share with sociology, social psychology, anthropology and perhaps economics, and it is eminently worthwhile. But there is no primacy of place for culture which is specifically digital.[3]

Humanities computing is likewise not the cultivation of an applied branch of the humanities, a goal urged upon us by a recent Dutch policy document (Bijker and Peperkamp, 2002). To be fair, let us note that Bijker and Peperkamp urge a redirection of the goals of humanities scholarship in general, not those of humanities computing in particular. Like most modern scientists and scholars, we are eager to pursue applications when these naturally arise,[4] but we seek a scientific direction that is worthwhile and independent of utility in applications.

De Smedt (2002, p. 90) is convincing in his plea to eschew pedagogical goals as primary in defining humanities computing curricula, pointing rather to the need to derive pedagogical goals from the scholarship of the field. De Smedt has it exactly right: university-level education needs to introduce students to research, to that terrible, exciting step in learning where one attempts to discover something new. We need to inform instructional plans via scholarly goals.[5]

COMMON Methodology?

McCarty urges that humanities computing be taken to embrace common humanities methodologies (McCarty 1998). McCarty is inspiring and wise about many aspects of humanities computing, and his views are not incongruent with ours, in particular his sympathetic recognition of the legitimacy of a colleague's request for reference to results (McCarty 1998, 'Introduction'). But McCarty overreaches:

> Humanities computing is an academic field concerned with the application of computing tools to arts and humanities data [...] It studies the sociology and epistemology of knowledge as these are

3 We should expect digital culture to be an attractive object of study, if only because it requires no digitalization of source material prior to analysis, but it is challenging to find examples of success where computational studies have contributed to the understanding of digital culture.

4 We have in fact been quite active in seeking applications in Groningen. See Nerbonne *et al.* (1998) and Gaustad and Bouma (2002) for examples of applications in computer-assisted language learning and text classification in service of language instruction and automatic email processing, respectively.

5 To avoid misunderstanding, let me emphasize that we have supported several efforts to harness the computer in university-level pedagogy. We have been involved in projects to improve instruction, e.g. in knowledge-based techniques (tcw.ppsw.rug.nl/nl/bok/), and in computer-mediated communication (www.let.rug.nl/cmc/). We find this worthwhile and interesting. But when we consider what needs to be done in humanities computing, a premature emphasis on pedagogy puts the cart before the horse.

affected by computing as well as the fundamental cognitive problem of how we know what we know.

McCarty invites *ad hominem* speculation about his views, noting of courses in humanities computing: ''The participants' disciplinary diversity has taught me that the only possible academic subject is the methodologies we have in common'' (McCarty 1999, II.A). McCarty may even be right about this, but we resist the conclusion that we need to find defining properties in methodology or foundational studies, preferring rather to emphasize the humanities and its research questions—even at the risk of having nothing more in common than what the larger humanities disciplines have in common.

McCarty will certainly agree that we need to pursue solid research results, but while we have emphasized that refection on methodology and foundation follows naturally in conducting research, McCarty stresses the search for common methodology. We suggest that the ambition to determine a 'methodological commons' for humanities computing—albeit a legitimate scholarly aim—is ill-directed vis-à-vis the challenges humanities computing now faces. If we focus on research results, we will naturally turn to methodological and foundational questions as part of the scholarly process, and it will not matter whether we share these throughout humanities computing. Now, we need results more than we need reflection.

6 Conclusions and Prospects

Hugo Brandt Corstius, a Dutch colleague in computational linguistics, once defined that field, freely following Clausewitz, as *die Fortsetzung der Sprachwissenschaft mit anderen Mitteln.* We suggest distributing Brandt Corstius's recipe across the humanities disciplines, pursuing humanities research with computational means. Even though we learn from each other regularly, we do not suggest focusing on common or essential elements of humanities computing, and it would handicap the program to be limited in this way. Perhaps some elements will eventually be common, but they first need to be tested and proven in the various subfields. Focusing on commonalities too early would limit our understanding of humanities computing unnecessarily.

The view of humanities computing as a federation of disciplines carries a burden with it. Just as in other loosely interdisciplinary endeavors, results in humanities computing really need to be reviewed twice, once by the disciplinary experts, i.e. the linguists, historians or archaeologists, and once again by the computing experts. Worse, the well-known fallacy of composition lurks around every corner—it is conceivable that a given piece of work makes good disciplinary sense *and* represents good computing, but that the specific combination is problematic. To guard against the last danger, we need more humanities computing experts—people who understand how computational means are applied to humanities questions.

6.1 *The future*

For humanities computing to survive as an academic field, it needs to prove its value to our peers in humanities. We can do this best if we provide answers to research questions they are asking. The answers need to be convincing, they need to withstand critical scrutiny, and they need to generalize to new areas of investigation.

Naturally, the dynamic of investigation will not stop there. Instead, we will naturally be challenged to defend our claims, to analyze our methods, and to reflect on our successes and on our limitations. The process of reflection is built into scientific discourse. In any case, we need not seek this step separately—it will arise naturally. In this way we should hope to resume the reflection, which some see as the primary task, but in a later phase of research, and informed by more and more substantial contributions. Patience will be repaid by a more solid base from which to reflect.

The immediate and very exciting task is to research humanities questions computationally, seeking opportunities for distinctive contributions in particular in areas where abundant data will allow us fresh perspectives and new opportunities for results.

Acknowledgements

The views expressed here would never have arisen without a long dialogue with my colleagues in Humanities Computing (*alfa*-informatica) in Groningen. I am especially indebted to my fellow staff members Leonie Bosveld, Gosse Bouma, Harry Gaylord, Henny Klein, Elwin Koster, Gertjan van Noord, and George Welling, but also to the postdocs who have worked in Groningen over the past years, namely, Jan Daciuk, James Hammerton, Wilbert Heeringa, Rob Koeling, Rob Malouf, Mark-Jan Nederhof, Miles Osborne, and Jennifer Spenader, and to the graduate students, who keep showing us where we are wrong, in particular Leonoor van der Beek, Jacky Benavides, Marcus Bergman, Tamás Biró, Francisco Borges, Gerlof Bouma, Ismail Fahmi, Tanja Gaustad, Stasinos Konstantopoulos, Tony Mullen, Robbert Prins, Susanne Schoof, Christine Siedle, Erik Tjong Kim Sang, Hideotoshi Shiraishi, Ivo Stoianov and Begoña Villada.

References

Bijker, Wybe and Peperkamp, Ben (2002). *Geëngageerde geesteswetenschappen: perspectieven op cultuurveranderingen in een digitaliserend tijdperk*. Den Haag: Adviesraad voor het Wetenschaps- en Technologiebeleid (AWT).

Boonstra, Onno, Breure, Leen, and Doorn, Peter (2004). *Past, Present and Future of Historical Information Science.* Amsterdam: NIWI-KNAW. Also publ. as *Historical Social Research/Historische Sozialfoschung* **29**(2).

Bouma, Gosse and Kloosterman, Geert (2002). Querying Dependency Treebanks in XML. In *Proceedings of the Third international conference on Language Resources and Evaluation (LREC)*. Gran Canaria, pp. 1686–1691.

Chambers, J. K. and Trudgill, Peter (1998, [1980]). *Dialectology*. Pages quoted from the 1st edn. Cambridge: Cambridge University Press.

De Smedt, Koenraad (2002). Some reflections on studies in humanities computing. *Literary and Linguistic Computing*, **17**(1): 89–101.

Gaustad, Tanja and Bouma, Gosse (2002). Accurate stemming of Dutch for text classification. In Theune, Mariët, Nijholt, Anton, and Hondrop, Hendri (eds), *Computational Linguistics in the Netherlands (CLIN)2001*, Vol. 12. Amsterdam and New York: Rodopi, pp. 104–117. Language and Computers: Studies in Practical Linguistics 45.

Gooskens, Charlotte and Heeringa, Wilbert (2004). Perceptual evaluation of Levenshtein dialect distance measurements using Norwegian dialect data. *Language Variation and Change*, **16**(3): 189–207.

Gooskens, Charlotte (2004). Norwegian dialect distances geographically explained. In Gunnarson, Britt-Louise, Bergstrom, Lena, Eklund, Gerd, Fridella, Staffan, Hansen, Lise H., Karstadt, Angela, Nordberg, Bengt, Sundgren, Eva, and Thelander, Mats (eds), *Language Variation in Europe. Papers from the Second International Conference on Language Variation in Europe ICLAVE 2, , 12–14 June 2003*. Uppsala, Sweden: Uppsala University, pp. 195–206.

Gooskens, Charlotte, N.d. Traveling time as a predictor of linguistic distance. *Dialectologia et Geolinguistica.* Forthcoming.

Heeringa, Wilbert (2004). *Measuring Dialect Pronunciation Differences using Levenshtein Distance*. Ph.D. thesis, Rijksuniversiteit Groningen.

Heeringa, Wilbert and Nerbonne, John (2000). Change, convergence and divergence among Dutch and Frisian. In Boersma, P., Breuker, Ph. H., Jansma, L. G., and van der Vaart, J. (eds), *Philologia Frisica Anno 1999. Lezingen fan it fyftjinde Frysk filologekongres*. Ljouwert: Fryske Akademy, pp. 88–109.

Heeringa, Wilbert and Nerbonne, John (2002). Dialect areas and dialect continua. *Language Variation and Change*, **13**: 375–400.

Heeringa, Wilbert and Gooskens, Charlotte (2003). Norwegian dialect examined perceptually and acoustically. *Computers and the Humanities*, **37**(3): 293–315.

Heeringa, Wilbert, Nerbonne, John, Niebaum, Hermann, and Nieuweboer, Rogier (2000). Measuring Dutch-German contact in and around Bentheim. In Gilbers, Dicky, Nerbonne, John, and Schaeken, Jos (eds), *Languages in Contact*. Amsterdam-Atlanta: Rodopi, pp. 145–156.

Heeringa, Wilbert, Nerbonne, John, and Kleiweg, Peter (2002). Validating Dialect Comparison Methods. In Gaul, Wolfgang and Ritter, Gerd (eds), *Proceedings of the 24th Annual Meeting of the Gesellschaft fur Klassifikation*. Heidelberg: Springer, pp. 445–452.

Hopcroft, John, E. (1987). Computer science: the emergence of a discipline. *Communications of the Association for Computing Machinery*, **30**(3): 198–202. Turing Award Lecture.

Kleiweg, Peter, Nerbonne, John, and Bosveld, Leonie (2004). Geographic Projection of Cluster Composites. In Blackwell, Alan, Marriott, Kim, and Shimojima, Atsushi (eds), *Diagrammatic Representation and Inference. Third International Conference, Diagrams 2004, Cambridge, UK,* *Number 2980 in 'Lecture Notes in Artificial Intelligence'*. Berlin: Springer, pp. 392–394.

Koster, Elwin (2001). *Stadsmorfologie: Een proeve van vormgericht onderzoek ten behoeve van stedenbouwhistorisch onderzoek*. Ph.D. thesis, Rijksuniversiteit Groningen.

Kretzschmar, William, A., Jr, McDavid, Virginia G., Lerud, Theodore K., and Johnson, Ellen (ed) (1994). *Handbook of the Linguistic Atlas of the South Atlantic States*. Chicago: The University of Chicago Press.

Malouf, Robert and van Noord, Gertjan (2004). Wide coverage parsing with stochastic attribute value grammars. In *IJCNLP-04 Workshop Beyond Shallow Analyses—Formalisms and Statistical Modeling for Deep Analyses.* Hainan: ACL.

McCarty, Willard (1998). *What is Humanities Computing? Toward a Definition of the Field*. Available at www.kcl.ac.uk/humanities/cch/wlni/.

Nerbonne, John (1998). *Elektronische Incunabelen*. Groningen: CLCG. Available at www.let.rug.nl/~nerbonne/oratie/.

Nerbonne, John and Kleiweg, Peter (2003). Lexical Variation in LAMSAS. *Computers and the Humanities*, **37**(3): 339–357. Special Iss. on Computational Methods in Dialectometry, ed. by Nerbonne, John and Kretzschmar, William Jr.

Nerbonne, John, Dokter, Duco, and Smit, Petra (1998). Morphological processing and computer-assisted language learning. *Computer-Assisted Language Learning*, **11**(5): 421–37. Special issue on French edited by Michael Zock.

Nerbonne, John, Heeringa, Wilbert, and Kleiweg, Peter (1999). Edit distance and dialect proximity. In Sankoff, David and Kruskal, Joseph (eds), *Time Warps, String Edits and Macromolecules: The Theory and Practice of Sequence Comparison*, 2nd edn. Stanford, CA: CSLI, pp. v–xv.

van der Beek, Leonoor, Bouma, Gosse, Malouf, Robert, and van Noord, Gertjan (2002). The Alpino Dependency Treebank. In Theune, Mariët, Nijholt, Anton, and Hondorp, Hendri (eds), *Computational Linguistics in the Netherlands. CLIN 2001. Selected papers from the Twelfth CLIN Meeting*. Amsterdam: Rodopi, pp. 8–22.

van Gemert, Ilse (2002). *Het geografisch verklaren van dialectafstanden met een geografisch informatiesysteem (GIS)*. Master's thesis, Rijksuniversiteit Groningen. Available at www.let.rug.nl/alfa/scripties.html.

van Noord, Gertjan (2004). Error mining for wide-coverage grammar engineering. In *Proc. of 42nd Meeting of the Association for Computational Linguistics*. Barcelona: ACL, pp. 446–453.

Welling, George (1998). *The Prize of Neutrality: Trade Relations between Amsterdam and North America 1771–1871. A Study in Computational History* Ph.D. thesis, Rijksuniversiteit Groningen.