Projecting the Impact of Computers on Work in 2030

Stuart W. Elliott[1]
National Research Council

Prepared for the
Workshop on Research Evidence Related to Future Skill Demands
Center for Education
National Research Council
May 31-June 1, 2007

This paper describes an approach to projecting changes in skill demand that will result from increases in computer abilities over the next few decades. Rather than the typical approach of extrapolating from recent changes in skill demand—which focuses on older computer techniques that have already been broadly applied—this approach uses current research in computer science to indicate what new computer abilities could be applied broadly in the future. These new computer abilities are viewed through the lens of the human abilities assessed by the O*NET database of occupational characteristics. A pilot effort in applying the approach suggests that computers may displace humans in occupations accounting for 60 percent of the current workforce by 2030. Such change would require broad increases in workforce skills. The pilot effort suggests it is feasible to analyze how current computer abilities align with occupational requirements as a way of understanding likely future changes in skill demand. The paper concludes by arguing that this general approach should be used to mount a serious and sustained effort to project and help prepare for the extensive changes in work that are likely to occur as computer abilities increase over the coming decades.

---

[1] This paper would not have been possible without the early and generous support of the Russell Sage Foundation and the encouragement of Eric Wanner, Dick Murnane, Dick Cyert, and Herb Simon.

This paper describes an approach to projecting changes in skill demand that will result from increases in computer abilities over the next few decades. The details of a projection that would result from applying the approach rigorously are not described since this has not yet been carried out. Instead, the paper reports the results of carrying out a pilot version of the approach and then argues for the importance of undertaking a more rigorous application of the approach in the years ahead.

Although there are a number of important policy issues that will be profoundly affected by increasing computer abilities and their impact on work, this paper is motivated particularly by questions about the implications for future skill demands for current educational reform. It is important to try to project future skill demand in general if there is a possibility of substantial future change in that demand that might in turn require substantial change in the education and training systems we use to develop workforce skills. There is no particular need to project minor mismatches between skill demand and supply that can be solved by minor shifts in the education and training young adults receive after they have left high school. Minor adjustment can be accomplished through responses to labor market signals that identify occupations where skilled workers are particularly needed. However, it is important to identify any substantial mismatches between skill demand and supply that would require substantial changes in K-12 preparation. Young adults who are substantially ill-prepared by their K-12 education for the labor market they face are unlikely to be able to respond adequately to the labor market signals they receive.

Given the time necessary to implement change in K-12 education and for that change to have an effect on skill supply, it would be helpful to be able to anticipate any large changes in skill demand that may come in several decades. The current K-12 education system is preparing children who will enter the workforce between 2010 and 2020. When we contemplate any possible shifts in skill demand that might necessitate substantial change in the K-12 education system, we need to consider the system's adequacy for preparing children who will enter the workforce between 2020 and 2030. Although that period may seem far off, the time required to motivate and implement any substantial change in the educational system effectively prevents us from substantially affecting the education of earlier cohorts of children.

The premise of the approach advocated in this paper is that 1) the most important driver in skill demand change over the next few decades will be the change in computer technology, and 2) the best way to understand the skill demand implications of computer technology is by viewing that technology through the lens of human skills. Improved technology will increase the range of work skills computers can perform competitively, which will then shift the demand for skills supplied by the human work force towards those skills that computers cannot yet perform competitively. Although we can look at the effects of computers on work skills in the recent past, it is important to realize that the effect of computers on future work skills is likely to be substantially different than their effect on present work skills because of the continuing increases in computer capabilities. As a result, the approach advocated in this paper is to look directly at the computer science research literature as a way of understanding what work skills computers are likely to supply in the coming decades.

The paper is organized into seven sections. Section 1 argues that there is a significant likelihood of a large change in skill demand over the next several decades that will not be identified by techniques that look for skill demand changes that have occurred in the recent past. Section 2 outlines the proposed approach for projecting the future impact of computers on work and describes the pilot version carried out in the remainder of the paper. Section 3 aggregates information about occupational ability requirements into a form that can be compared to the computer science research literature. Section 4 synthesizes research about current computer abilities in terms of the ability levels used to assess occupations. Section 5 applies the assessments of computer ability levels to project changes in employment across occupations in 2030. Section 6 discusses some of the important policy issues raised by the pilot projections of 60 percent occupational displacement by 2030. Section 7 recommends that the proposed approach be used as the basis for a serious and sustained effort to project the impact of computers on work in the years ahead.

1. The likelihood of substantial change in skill demand before 2030

There is broad recognition of the substantial changes in computer capabilities and applications that have occurred over the past two decades, with the widespread adoption of personal computers and the internet for a number of business and personal uses. In addition, there is a general appreciation for the exponential increases in computer processing power that have made these changes in computer capabilities and applications possible. However, there is not a general appreciation for the impact these changes could have in the coming decades.

The exponential increase in computational processing power—commonly referred to in terms of Moore’s law, which relates to the 18-month doubling time for the number of transistors on an integrated circuit—has been going on for over a century (Moravec, 1999, p. 57). The exponential trend increases the computational power available at a given price by an order of magnitude every five years. Moore’s law is an observation about the most recent stage of this trend, which has occurred in integrated circuit technology over the past several decades. However, the trend extends back to 1900, persisting through the evolution of mechanical, electromechanical, vacuum tube, and transistor technologies. There is regular discussion in the press about new computing technologies that may continue future improvements in computational processing power at the point when further improvements in integrated circuits begin to fall short (e.g., Chang, 2007). There is a general expectation in the engineering and business community that exponential increases in computational processing power will continue to take place over the next few decades.

The exponential trend in computer processing power means that inexpensive personal computers in a couple of decades will be able to perform tasks that are simply impractical for current personal computers to perform. Imagine, for example, some task X that humans are able to perform in a second and suppose that personal computers are able to perform that same task but that it takes them four orders of magnitude longer to perform it. As a result, the one-second human task becomes a three-hour task on today's computers. However, in two decades, the exponential increase in computing power will mean that the one-second human task is also a one-second computer task. And a decade after that, personal computers will be able to perform the task one hundred times in the time humans take to perform the task just once. Thus, over a period of three decades, computers are transformed from being completely impractical for performing task X to performing the task far more quickly and cheaply than humans are able to do.

Given current technology and the likely continuation of the exponential trend, personal computers are likely to approach the computing power of the human brain over the next several decades. An extrapolation from the computer processing necessary to reproduce the brain's low-level visual processing suggests that reproducing the processing power of a full human brain would require roughly 100 trillion instructions per second (Moravec, 1999, p. 54). For comparison, personal computers today provide on the order of 10 billion instructions per second, four orders of magnitude slower than the processing power necessary to reproduce the brain's function. As a result, there are a number of the brain's functions that current personal computers are simply too slow to be able to carry out effectively. However, a continuation of the past exponential increase in computer processing power would allow computers to close this gap by 2030.

The general exponential increase in computer processing power means that we should expect substantial increases in computer capabilities over any period of several decades, but the comparison between human and computer processing power suggests that we should expect a particularly critical set of changes over the next period of several decades. By 2030, it is plausible to think that personal computers will be able to perform a number of human work tasks faster and more cheaply than humans are able to do. If this change does indeed occur, it could cause a large restructuring of the types of occupations that humans perform—exactly the type of large structural change that we might want to prepare students for in advance. Depending on the nature of the occupational restructuring, it could prove difficult for many new labor market entrants to find work in 2030, unless they have a substantially different kind of K-12 education than they are currently likely to receive.

The existing approaches to measuring changes in skill demand will not be helpful in projecting this potential shift in skill demand because they look backwards rather than forwards. The existing approaches look at changes in the skill content within an occupation or changes in pay and number of jobs across occupations (CFE workshop refs). Of necessity, these approaches look at the effects of technology that has already diffused and been applied to the labor market. So the information we have from such approaches in 2007 tells us about changes that occurred in the late 1990s as technology diffused to the workplace that first became commercially available in the late 1980s and early 1990s and that was developed still earlier. However, in projecting potential changes to skill demand that might lead us to consider some changes in K-12 education, we need to know about technologies that may not be commercially available until the

2010s and may not be broadly applied until the 2020s. That is, we need to know about the impact of computer technology that is four-to-six orders of magnitude more powerful than the computer technology affecting the labor market of the late 1990s that is subject of study of the existing approaches to measuring changes in skill demand. Given the qualitative change in functionality that will be brought by four-to-six orders of magnitude of change, the findings of the existing approaches to measuring skill demand changes are simply inadequate for making the projections that we require.

2. An alternative approach to projecting changes in skill demand

As noted above, the approach advocated in this paper for projecting future skill demand changes is to look directly at the computer science research literature as a way of understanding what work skills computers are likely to supply in the coming decades. The basic premise is that computer abilities that are being actively investigated now are likely to be refined over the next decade and then diffuse broadly through the economy in the decade after that. Thus, the suggestion is that we can use the current research literature as a guide of computer abilities that are likely to be used broadly throughout the economy by 2030.

However, the computer science literature is large and it can be difficult to determine how to translate the various capabilities discussed into implications for future computer abilities in relation to human work. To bridge this gap, the approach used here begins with a preliminary step of characterizing the various abilities used in human work. These abilities are then clustered in a way that can be compared to the computer science literature.

The paper carries out an analysis of the impact on computer technology on future skills in the following three steps:

- First, the U.S. Department of Labor's database on occupational characteristics—O*NET—is used to define the set of human abilities that are broadly relevant to work, along with anchoring tasks for different difficulty levels for those abilities and ratings for occupations on each ability scale. The abilities are sorted into four groups that are used to characterize the skill requirements of the full range of current occupations.

- Second, recent articles in *AI Magazine*—a journal with broad coverage of research and innovative applications in artificial intelligence—are used to define computer capabilities in the four groups of abilities drawn from O*NET. The O*NET anchoring tasks are used to place computers on the ability scales used to rate occupations.

- Third, the ability levels of current research computer systems described in *AI Magazine* are then used as a rough guide to the computer capabilities that will be refined and broadly applied by 2030. These capabilities are used to identify occupations that are likely to experience substantial levels of automation, using the simplistic inference that occupations will be automated when computers become capable of performing the full range of required abilities. The mix of abilities in the occupations that are *not* projected to be automated are then used to define the distribution of skill demand for human workers.

The three-step analysis as carried out in the present paper is intended to serve only as a pilot demonstration of a more rigorous effort. Obviously, each of the three steps above involves gross approximations. The analysis could be substantially improved by a number of refinements:

- Analyzing the O*NET ability scales separately rather than in groups
- Incorporating information from other O*NET scales
- Drawing from multiple sources to describe current computer research
- Using analysis from computer science for insight about which computer skills may advance fastest in the near future
- Using analysis from cognitive psychology to define benchmarks for human performance levels
- Using ethnographic studies to investigate the plausibility of automation in individual occupations that the approach suggests will be automated
- Using occupational analysis to identify occupations that could be substantially restructured without computers being capable of performing the full range of capabilities
- Using economic analysis to project the shifting mix of occupations that will occur as some work becomes cheaper to perform through automation

Such refinements can be included in later efforts. The purpose of the present paper is merely to demonstrate that there is some insight to be gained about future skill demand from an attempt to understand what new computer capabilities may have a broad effect on work by 2030. The next three sections of the paper carry out the steps of the pilot effort outlined above.

3. Step one: The ability requirements of human work in the U.S. in 2005[1]

Over the past 15 years, the U.S. Department of Labor, in collaboration with several states, has developed a new database of occupational characteristics—O*NET—to replace the outdated Dictionary of Occupational Titles (Peterson et al., 1995, ch. 1). The

---

[1] If data were available, it would be useful to carry out this analysis globally.

database will eventually include information about most of the occupations that are classified within the Standard Occupational Classification (SOC).[2] Data are being collected from job incumbents and occupation experts in an ongoing process that is populating the database in a series of waves. These waves of data-collection are expected to continue after the database is fully populated in order to successively provide revised information about all occupations, including any new occupations that are added to the occupational taxonomy.

The latest version of the O*NET database—version 11.0—includes information on 798 occupations. For most of these occupations the ratings are based on recent information provided by job incumbents and analyzed by job experts, those for some occupations the ratings are still based on analyses of earlier data provided by job experts alone. The occupations included in version 11.0 cover roughly 95 percent of the current workforce.

For each occupation, the database includes information about 12 different sets of descriptors related to worker characteristics, requirements, and experience, along with information about 6 different sets of descriptors related to the tasks and activities of the occupations themselves. Within each one of these 18 sets, there are a number of individual scales or indicators. In total, the database includes information about several hundred different descriptors. In addition, through the SOC, O*NET can be linked to other occupational information, such as data from the U.S. Bureau of Labor Statistics (BLS) on employment.

Table 1 lists the 18 different sets of O*NET descriptors for workers and occupations. All of these descriptor sets provide potentially useful information about the worker capabilities and job requirements that computer technology would need to meet to be able to substantially automate an occupation. Not surprisingly, there are substantial overlaps across the different scales. For example, the Abilities set of descriptors has a scale for written comprehension, the Basic Skills set has a scale for reading comprehension, the Education set asks about education requirements related to English and language arts, the Basic Skills – Entry Requirement set asks about entry requirements related to reading comprehension, and the Generalized Work Activities set has one scale related to getting information and two related to communicating with others. Given the limited time available for this pilot effort, the current paper focuses on the Abilities set of descriptors. This set is defined as reflecting "enduring attributes of the individual that influence performance" and it appears to describe these attributes in a sufficiently general way to allow successful comparisons to corresponding computer abilities as they are described in the research literature.

Table 2 lists the set of 52 scales that make up the Abilities set of descriptors in O*NET. The 52 scales are grouped into 4 major clusters and 15 sub-clusters. An asterisk is placed next to the 22 scales that are used in the pilot analysis. The other scales are ignored because a cursory consideration suggests that technology already surpasses human abilities. In most cases—such as the four abilities related to strength or the ability related

---

[2] Most of the information on the O*NET program has been drawn from the project's website at http://www.onetcenter.org/.

to number facility—technology's current superiority seems to be relatively uncontroversial. However, at least two of the omitted scales may deserve a brief explanation:

- Fluency of ideas – although humans clearly still have superiority in the general area of reasoning, this particular ability specifically focuses on the sheer number of related ideas generated: "the *number* of ideas is important not their quality, correctness, or creativity." Given the ability of search engines to generate huge quantities of "related" information, it seems appropriate to give this scale to the computers.

- Speech recognition – although computers may not have quite reached human levels of speech recognition, most of the remaining limitations appear to be related to language understanding. For speech recognition per se, the earlier limitations of speech recognition systems—which required small vocabularies, word separation, and training to individual speakers—have been mostly overcome (ref).

In any case, given the likely correlation across the different ability scales and the rough approach to this pilot analysis, it's unlikely that the results will be significantly affected by the inclusion or exclusion of several ability scales.

For the pilot analysis, the 22 ability scales where humans are still superior are placed into the four groups shown in Table 3. These four groups—language, reasoning, vision, and movement—are closely related to the clusters and sub-clusters used in O*NET and have the benefit of grouping abilities that are likely to be addressed by similar projects in the computer science literature.[3]

For collecting information about the different ability scales from job incumbents and occupation experts, the O*NET questionnaires use 7-point scales with three anchoring tasks provided for each ability. The anchoring tasks were developed and tested during the development of O*NET, with examples chosen for each ability that respondents placed in relatively consistent positions on the scale (Peterson et al., 1995, p. 10-22). In most cases, the three anchoring tasks that were chosen are ones that respondents placed on average at levels 2, 4 and 6, respectively. However, in some cases, one or more of the anchors chosen was from a different level on the scale. Table 3 provides the low, medium, and high anchors for the 23 ability scales. The table notes cases where the anchor is placed on the ability scale at a position other than 2, 4 or 6.

To make the pilot task manageable, the remainder of the paper focuses on the four groups of abilities rather than on the 22 individual abilities that make up those groups. Aside from the appeal of making the effort doable, this rough approach to considering information about the abilities is justified by two plausible assumptions:

---

[3] Note that in the case of Visualization, it seemed more appropriate to group the ability with other reasoning abilities—because of its focus on higher level visual processing and reasoning—rather than with other vision abilities.

- The intuitions of respondents about the characteristics that make a task more or less difficult with respect to a single ability are likely to be similar across the different abilities in a group.

- The requirements for computers to reproduce different levels of the abilities are likely to be similar across different abilities in a group.

So, for example, within the language group, it is plausible that people have similar intuitions about the characteristics that make a written comprehension task more or less difficult as they have about the characteristics that make an oral expression task more or less difficult. And it is plausible that the requirements for computers that can read are likely to be similar to those for computers than can speak. In both cases—the evaluation of job tasks and the provision of computer capabilities—it is plausible that the different levels of difficulty and complexity are distinguished by such things as overall length, difficulty of the required vocabulary and grammatical structures, and familiarity and complexity of the knowledge being communicated.

The rated occupations in the O*NET database include average ratings for each ability scale for each occupation. Typically, these average ratings reflect the input from eight job incumbents who replied using the 7-point scale. The averages are reported to the tenths position and the standard error of the average is typically reported to be around 0.3.

To aggregate the separate ability scales within each of the four groups, the maximum rating across the separate scales is used. The argument for taking the maximum rather than the average is that full automation of that group of abilities would require computers to be capable of the full set of abilities in the group. Thus a job requiring a low level of written comprehension but a high level of oral expression is rated as requiring a high level of language ability. After taking the maximum, the rating for each ability group for each occupation is rounded to the nearest whole number.

Table 4 shows the distribution of employment across the seven ability levels for the four different ability groups. Employment for each occupation is derived from 2004 figures from the BLS.[4] Not surprisingly, the employment distribution within each of the four ability groups reflects a rough normal distribution. For each scale, the bulk of employment is clustered over three levels. For language and reasoning most employment is clustered on levels 3-5, whereas for vision and movement most employment is clustered on levels 2-4. Only a small portion of employment requires abilities at level 6, the level typically associated with the high anchor on the ability scales.

---

[4] Employment data were downloaded from the BLS at http://www.bls.gov/oes/oes_dl.htm. Employment in the occupational groups omitted from the current version of O*NET—representing 5 percent of employment—is ignored and the distributions are calculated from the included occupations alone. Effectively this treats the omitted occupational groups as requiring a similar distribution of skills as the included occupations. Since many of the missing occupational groups are the “other” occupational categories within the broader occupational groups—such as 13-2099 Financial specialists, all other—this is probably a reasonable assumption.

To provide a sense of the ability level ratings for specific occupations, Table 5 provides average ratings for the four ability groups for the 93 minor groups of occupations in the SOC.[5] The ability averages across individual occupations are weighted by employment for each occupation in the minor group. The table also provides the overall percentage of employment represented by each minor group. In general, the table shows the kinds of contrasts across occupations that one would expect to find, such as higher requirements for language and reasoning ability among managers, engineers, and scientists, and higher requirements for vision and movement ability among doctors, protective service workers, and motor vehicle operators.

To provide a rough picture of the distribution of employment across the full set of abilities, Table 6 aggregates the four ability groups into two, combining language and reasoning abilities together into a single group and similarly combining vision and movement abilities. The aggregation is performed by taking the maximum across the two groups. The table shows the percentage of current employment in occupations requiring each combination of language-reasoning (LR) abilities and vision-movement (VM) abilities.[6]

It is interesting to note the suggestion from Table 6 that substantial levels of automation of current employment is likely to require progress on both LR and VM abilities. There are not many current occupations where the LR or VM abilities are at level 1. We may not want to take these ability estimates at face value, since they reflect abilities used in occupations that are currently carried out by humans who generally possess the full range of abilities and so in general there’s no need to try to design a job to avoid the need for a particular portion of the abilities that humans generally have. However, if we do take the ability estimates at face value, they suggest that substantial levels of automation may not occur until computers are capable of a fairly full range of the abilities humans generally possess.

4. Step two: Computer abilities in the current research literature

For the second step of the pilot analysis, recent articles from *AI Magazine* were used to provide concrete examples of the ability levels of current computer systems in the research literature in computer science. *AI Magazine* is a quarterly journal with broad coverage of areas in artificial intelligence that is published by the Association for the Advancement of Artificial Intelligence. The articles in this journal are technical without being so detailed as to be difficult to understand for a non-specialist. A review of the

---

[5] The entire pilot analysis excludes occupations in the major group of the Military Specific Occupations and its three minor groups because they are omitted from the current version of the O*NET database.

[6] The employment percentages in Table 6 don’t correspond exactly to those shown in Table 5 because the ability combinations of the individual occupations that are aggregated in Table 6 are sometimes different than the average ability combinations of the minor occupational groups. In addition, the Table 5 employment percentages reflect the full employment in the minor occupational groups, effectively including the employment in the small set of occupations that are not yet included in O*NET.

titles and abstracts for the articles published between 2003 and 2006 produced a set of articles that collectively addresses each of the four ability groups

The objective for step two was to find several articles within each of the four ability areas that give concrete examples of the ability levels of current computer systems that can be compared to the anchoring tasks of the O*NET ability questionnaires. Therefore, in selecting articles, it was important to look for those that provide concrete examples of tasks that the research systems can perform. Many of the articles in the computer science literature focus on promising new techniques and do not provide good examples of the capabilities of systems using those techniques. In addition, *AI Magazine* includes a number articles describing recent conferences, influential researchers, as well as retrospective and prospective pieces about the field; in almost all cases these articles are too general to provide examples of concrete skills that could be compared to the anchoring tasks of the O*NET ability questionnaires.

The four-year review produced 12 articles that describe computer systems at the appropriate level of detail. The computer abilities described in each of these articles are summarized in Table 7, with specific attention to noting limitations in performance that are included in the articles. Statements about the long-range goals of the different research programs are kept to a minimum since they do not describe current capabilities.

The remainder of this section discusses the four ability group in turn. The section for each group begins by reviewing the anchoring tasks on the O*NET ability questionnaires and the contrasts among the different ability levels. This is followed by a synthetic description of the computer abilities drawn from the 12 summaries in Table 7 and then a tentative suggestion about the O*NET ability level that corresponds to these abilities. As noted above, the argument is that it is plausible to think that such research capabilities will be refined over the coming decade and applied in the economy in the decade after that, with the result that computer abilities at the designated ability level could be broadly applied throughout the economy by 2030.

a. Computer language abilities

There are four O*NET scales in the language group of Table 3a, related to listening, reading, speaking, and writing. The meaning of the scales is straightforward, with two related to understanding, two related to expression, two related to oral language and two related to written language.

The O*NET technical manual for the development of the content model (Peterson et al., 1995, Appendix G) provides general descriptions at the low and high ends of the ability scales that are intended to capture the contrasts in ability difficulty represented on each scale. On all four scales, the basic contrast is between language involving simple and complex ideas. For some of the scales, there are elaborations of this contrast, with the low end involving short sentences and common words and the high end involving detailed sentences, clear organization, and unusual words.

At the low end of the anchoring tasks provided in Table 3a, the language use can involve a single word or phrase, such as a street sign, or a few sentences, such as in a television commercial or a brief request to a customer service representative. At the high end, the language anchoring tasks involve extended and advanced material, such as the understanding or production of a lecture or book in a technical area. In between these extremes, the medium level of anchoring tasks involves language use corresponding to a page or two of material, such as an apartment lease, a recommendation letter, or a set of multi-step instructions.

The systems summarized in Table 7 include six that involve some aspect of language use (*abeijk*).[7] The systems span the four language scales of O*NET, with one involving written comprehension only (*a*), two involving written expression only (*ei*), two involving oral comprehension and expression (*jk*), and one involving all four abilities (*b*).

The systems involve a diverse range of tasks that includes detecting problematic text in an insurance application (*a*), providing customer service for both sales and repairs (*b*), explaining the answers to chemistry questions in an AP-style test (*e*), describing the movement of cars in a video of a traffic intersection (*i*), asking for help in finding the registration booth at a conference (*j*), giving a conference talk that includes questions from the audience (*j*), and role-playing with students in a training simulation about how a military officer should handle a car accident with a civilian (*k*). Two of the systems have been applied commercially (*ab*).

One notable point about the six systems involving language use is that they all involve integration with one or more of the other three ability groups. In all cases the integration involves reasoning, in two it involves vision (*ij*) and in one it involves movement (*j*). The integration of different ability groups in these systems indicates that the understanding of appropriate techniques has progressed far enough that it is possible for a single research group to combine techniques in several different ability groups rather than focusing on the challenges of a single ability group alone. This has important implications for the capacity of computers to take over significant portions of human occupations, since the O*NET analysis indicates that most jobs require significant levels of all four ability groups (Table 5).

In terms of the length and complexity of the language these systems can work with, they seem to be closest to the middle level of anchoring tasks in Table 3: they are beyond language use at the word or couple-sentence level but fall far short of language use at the lecture or book level. The only system that suggests language use at the high level is the system that gave a conference talk (*j*), but while that system monitored its progress through the concepts it needed to discuss—so that it could flexibly answer questions without repeating an in-depth description that had already been given—the system did not independently develop the overall structure of its talk.

In addition to length and complexity corresponding to a page or two, the medium-level anchoring tasks for expression—giving directions to a lost motorist and writing a job

---

[7] Italicized letters in parentheses refer to the 12 different systems summarized in Table 7.

recommendation—both appear to involve some ability to adjust language to the needs of the person who is being communicated with and the requirements of the situation. Several systems exhibited this kind of sensitivity, including the ability of the conference talk system to monitor its points and not repeat them (*j*) and the ability of the training simulation to reason about emotion in order to choose appropriate language and understand imprecise language (*i*).

Other details provided about the systems also seem to be consistent with performance at the middle ability level for language. For example, the training simulation has a speech recognition vocabulary of hundreds of words, a vocabulary for language generation of a thousand words, and an understanding of 40 tasks and 150 properties of the world (*i*). The system replies verbally in three seconds to a short verbal input. These are systems that can use language related to a relatively defined set of material. One of the systems specifically provides a backup approach for dealing with people who press beyond the limits of the material that the system can handle adequately (*b*).

For a projection of computer language ability to 2030, the systems in Table 7 suggest it is plausible to expect the widespread application of computers using language at level 4, the medium level of difficulty of the O*NET scale.

Going beyond the capabilities of the systems described in Table 7, it is important to note the work of the Cyc research group, a controversial and ambitious project that has been underway for over 20 years (Lenat, 1995). The project is encoding the full range of common sense knowledge humans have about the world, with the objective of allowing computers to fully understand and use language in the way we do. One aspect of this objective is the capability for computers to begin to add to their own knowledge by reading, a step that would allow the rapid expansion of knowledge encoded in the system. The system currently has 3 million facts and rules about hundreds of thousands of concepts, providing far more comprehensive coverage of language and concepts than the systems described in Table 7. Until a couple years ago, all of the knowledge in the Cyc system had been entered “manually” by humans. The system is now sufficiently advanced that the group is working on ways for it to begin to add to its own knowledge base on its own (Shah et al., 2006). The project’s leader has stated optimistically that the system might be on track to reach human-level language and reasoning by 2020 (Lenat, 2006). This is a project that should be tracked closely because it has the potential for substantially increasing the level of computer language ability projected for 2030.

b. Computer reasoning abilities

There are six O*NET scales in the reasoning group of Table 3b. These involve recognizing that a problem exists, applying general rules to solve a problem, and developing new rules or conclusions.

As with the language abilities, the basic contrast between the low and high ends of the reasoning scales is between reasoning involving simple or common ideas on the one hand and complex or unusual ideas on the other (Peterson et al. 1995, Appendix G). There is a separate scale for “originality” that specifically focuses on the ability to develop "unusual," "clever," or "creative" ideas, but it is not clear that this has been defined in a way that clearly distinguishes it from the higher ends of the reasoning scales involving applying or developing rules where the result may well be novel. Note also that there are separate scales for mathematical reasoning and spatial reasoning, but it is not clear that these have been defined in a way that clearly distinguishes them from the more general scales of deductive and inductive reasoning.

At the low end of the anchoring tasks provided in Table 3b, the reasoning involves simple one- or two-step inferences, such as recognizing that a lamp won't work if it's unplugged, realizing that a stalled car can coast downhill, or calculating the cost of 10 oranges from their price. In many cases, the low end reasoning tasks are ones that would be classified as "common sense." At the high end, the reasoning involves problems with many features and requiring many inference steps, such as diagnosing a disease, designing an aircraft wing, simulating a spacecraft landing, or planning a chess move. The high end anchoring tasks are all drawn from domains that are generally recognized for requiring a high level of specialized knowledge and expertise, such as engineering, medicine, and chess. The medium level tasks seem to involve a mix of common sense reasoning and more specialized expertise, such as calculating profits, identifying a crime suspect, or following a diagram to assemble a cabinet.

The systems summarized in Table 7 include eight that involve some aspect of reasoning that relates to the six O*NET scales in the reasoning group (*abcdefgk*). The systems span the six reasoning scales, with two involving the recognition that a problem exists (*ag*), four involving the application of general rules to solve a problem (*abek*), two involving the development of novel rules or conclusions (*cf*), two involving mathematical reasoning (*be*), and two involving spatial reasoning (*df*). As noted in the previous section on language, a number of these systems integrate reasoning and language abilities (*abek*).

The tasks addressed by these systems include making underwriting decisions about long-term care insurance (*a*), providing customer service for both sales and repairs (*b*), developing new hypotheses about good conditions for growing crystals and for recovering from medical disability (*c*), providing useful analogies for solving problems in qualitative physics and in military tactical games (*d*), providing answers and explanations to chemistry questions in an AP-style test (*e*), developing novel atomic models for electron density maps of proteins (*f*), identifying patterns of potentially suspicious facts that could indicate a terrorist plan (*g*), and role-playing with students in a training simulation about how a military officer should handle a car accident with a civilian (*k*). Four of the systems have been applied commercially (*abcf*).

One of the striking aspects of computer systems in the area of reasoning is their ability to produce high levels of performance. For example, the systems in Table 7 make insurance underwriting decisions about easy cases and provide guidance to underwriters about more difficult ones (*a*), produce novel hypotheses about growing crystals that are sufficiently promising to merit further investigation (*c*), achieve scores on an AP chemistry exam comparable to the mean score for advanced high school students (*e*), and produce initial atomic models for proteins that substantially reduce the time for experts to develop refined models. This evidence of performance suggests reasoning at the higher level on the O*NET scales, though somewhat below the high level of the anchoring tasks since it seems the human experts are generally able to outperform the computers. Thus the systems in Table 7 suggest that computer reasoning ability might be placed at level 5 on the O*NET scale.

It is important to note, however, that the common sense reasoning that is typical of the lower levels of the O*NET reasoning scales has been surprisingly difficult for computers to perform in many cases. The explanation seems to be that the knowledge required for common sense is very large but is easily overlooked because everyone shares it. The reasoning systems in Table 7 rely on knowledge encoded in hundreds or thousands of rules, properties, or cases (*ck*), whereas the necessary knowledge base for common sense appears to involve millions of elements (Lenat, 1995). For specialized areas of reasoning, such as those that appear in the high level anchoring tasks, we are accustomed to providing humans with specialized knowledge and so it seems appropriate that computers also need specialized knowledge to perform those tasks. However, the success of computers on these specialized tasks does not mean we can then assume—as we can with most adults—that they are also capable of performing "simpler" tasks that rely on common sense. For computers to be able to perform these common sense tasks, they still need to be provided with the requisite knowledge. This has proven to be a very large undertaking, but it appears that the Cyc project has now made the necessary initial investment in codifying the knowledge for common sense to allow computers to perform the reasoning tasks at the lower levels of the O*NET reasoning scales. Without this investment, we might have wanted to conclude that computers were capable of level 5 reasoning tasks but perhaps not tasks at levels 2, 3 or 4. However, given that this investment has now been made, projecting computer reasoning abilities at level 5 for 2030 seems appropriate.

c. Computer vision abilities

There are four O*NET scales in the vision group of Table 3c. These involve recognizing objects and different features of those objects, including their position in space.

The basic contrast in the level of difficult between the low and high level anchoring tasks varies somewhat across the different scales, relating to the complexity of the object that is identified, the complexity of the background in which the object is recognized, or the complexity of features that need to evaluated.[8] This complexity can make an object hard to identify—because its own visual pattern is difficult or because other objects occlude or are similar to it—or can make it necessary to focus on precise or subtle features. In two cases, the scales specifically mention speed, but expected increases in processing power should make speed irrelevant by 2030 for any visual ability that is possible with current computers.

The low end anchoring tasks involve objects that are visually relatively simple and occur in uncluttered environments, such as identifying the zip code on a letter or judging the

---

[8] The two "closure" scales refer to perception more generally and the low end anchoring tasks are auditory rather than visual.

relative distances of cars in traffic. The high end anchoring tasks involve difficult patterns and distracting environments, such as identifying camouflaged tanks from a high-speed plane, inspecting parts for defects on a fast-moving assembly line, or identifying the precise location for throwing a football. In between these, the middle level uses example tasks such as making sense out of strange handwriting or looking for a golf ball in the rough.

The systems summarized in Table 7 include four that involve vision (*hijl*). All four of the systems involve object identification, ranging from balls and robots in simplified environments (*h*), to cars, roads, registration booths, elevators, hallways, and people in more naturalistic environments (*ijl*). All of the systems also involve recognizing features of the identified objects, particularly their location and movement. The tasks of the systems include locating a soccer ball and other soccer players (*h*), identifying cars and their movements in a video of a traffic intersection (*i*), finding the registration booth, an elevator and several rooms at a conference (*j*), and identifying drivable surfaces and obstacles for an autonomous car (*l*). All four systems involve integration, in three cases with movement (*hjl*), in one case with language (*i*), and in one case with language and reasoning (*j*). None of the systems has been deployed commercially.

None of the systems involve vision abilities at the high end of the anchoring tasks where the patterns are very difficult and the environment is highly confusing. The visual problem of one of the low-level anchoring tasks—that of locating the position of other cars in traffic in order to be able to merge while driving—is similar to the problem that these systems are trying to solve. None of the systems involve difficult patterns, background visual distraction, or subtle distance detection that would correspond to even the medium level of the vision anchoring tasks on the O*NET scales. These comparison suggest that current computer vision abilities should be placed at level 2, corresponding to the low level on the O*NET scales.

It is important to realize, however, that three of the four articles in *AI Magazine* describe vision systems that are part of an integrated robotic system. As a result of their use in a robot that needs to respond relatively quickly to its environment, the vision systems are forced to take processing shortcuts that keep them within the limits of currently available computing power. For example, the computer vision system driving the autonomous car relies on initial input from a laser guidance system to identify drivable surfaces—rather than on direct analysis of visual input—in order to allow the system to operate fast enough. And even with this shortcut, the limits on the vision system prevent the vehicle from driving faster than 38 mph (*l*).

To obtain a more accurate assessment of current computer vision abilities for projecting future skills, it may be preferable to focus specifically on systems that are not being used in integrated robotic systems and therefore have the luxury of operating slowly. Of course, greater processing power will allow such systems to be used in integrated robotic systems in the next decade or two. For example, Sebe and Lew (2003) describe a number of computer vision experiments. In one set, they describe the relative performance of several different approaches for matching stereo images for depth perception, using a variety of complicated naturalistic images. In another set of experiments, they describe the performance of a system for recognizing emotions from facial expressions. These examples suggest abilities that involve greater visual complexity and potential confusion than the systems in Table 7, perhaps closer to the medium level anchoring tasks. Given this additional information, it seems more appropriate to place current computer vision abilities at level 3, between the low and medium levels on the O*NET scales.

d. Computer movement abilities

There are eight O*NET scales in the movement group of Table 3d. These involve spatial orientation, coordination of various body parts, control of a movement or a piece of equipment, and overall body equilibrium.

The basic contrast in the level of difficulty between the low and high ends of the movement scales is specified in somewhat different ways across the different types of control. For spatial orientation, the difference in difficulty appears in whether the environment is unchanging or changing, with a changing environment requiring repeated adjustments to one's sense of orientation. For the four coordination scales, the difference in difficulty appears to be between rough or precise levels of movement and the number of body parts that need to be coordinated, as well as the speed of the coordination required. For the two control scales, the difference in difficulty appears in whether adjustments are required a few times or repeatedly and whether they are rough or precise. For the equilibrium scale, the difference in difficulty appears in whether balance is maintained against one or many forces of instability, with the latter presumably requiring more adjustments. Generalizing over these four different types of scales, movement difficulty seems to involve adjustments that are more frequent and more precise and that involve more body parts.

At the low end of the movement scales in Table 3d, the anchoring tasks require only a few steps and these can be fairly rough, such as locating a room from a map, putting coins in a parking meter, adjusting a light with a dimmer switch, or balancing on a ladder. At the high end, the anchoring tasks require repeated and precise adjustments, such as navigating on the ocean using the sun and stars, assembling a watch, drilling a tooth, or walking on beams in a high-rise construction site. In between these extremes, the medium level uses such example tasks as walking through a darkened room without hitting anything, operating a forklift in a warehouse, keeping up with a car that changes speed, and walking on ice across a pond.

The systems summarized in Table 7 include three that involve movement (*hjl*). All three systems involve spatial orientation, coordination, and control. Only one of the systems involves body equilibrium (*h*). The tasks of the systems include walking (*h*), kicking a ball (*h*), passing a ball between two robots (*h*), moving down a hallway (*j*), following a map to locate a meeting room in a hotel (*j*), using an elevator (*j*), and driving a car in the desert (*l*). All three systems involve integration with vision and one involves integration with language (*j*). None of the systems is being used commercially.

None of the movement abilities of the three systems comes close to the high-level anchoring tasks in Table 3. One of the systems carries out the low-level spatial orientation task of locating a room from a map (*j*). The driving system moves through an extensive environment and stays oriented within a changing road surface in that environment (*l*). All three systems involve coordination of a series of movements over an extended period of time. These movements are not very precise, although the driving system coordinates its steering and driving speed to stay on the road and avoid obstacles (*l*). The control abilities of the driving system—which is able to adjust its speed to changing road conditions—seems to be above the low level anchoring tasks which involve a single adjustment to a smooth change, but do not seem to be quite at the medium level of precision indicated by such tasks as operating a forklift or adjusting to the speed of another vehicle. Similarly, the equilibrium of the systems that walk and kick a ball seems to be above the low level task of balancing on a ladder but below the medium level task of walking on ice (*h*). Overall these comparisons suggest that the movement ability of these computer systems is roughly at level 3.

It is worth considering briefly the two future competitions that are described in Table 7. Although the future competitions obviously do not have direct implications for current computer abilities, they are likely to have been designed with an eye towards feasible developments that build on current movement capabilities. First, the DARPA competition that produced the independent driving system is turning next to a competition that requires vehicles that can drive in an urban setting, responding appropriately to other vehicles, as well as to traffic signs and regulations (*l*). The first trial for the new competition will be held later this year. This level of movement ability seems still between the low and medium levels of the O*NET scale, but suggests developments over the next decade or so towards abilities, such as operating a forklift, that would place computers solidly at the medium level. Second, the systems for walking and kicking balls have been developed in the context of RoboCup, an annual robot soccer competition that has the objective of fielding a team of robots by 2050 that can play soccer according to the normal rules and can defeat the human world cup champions (*h*). This level of movement ability is at the high level of the O*NET scale and would clearly suggest substantial additional improvement over the coming four or five decades. In both cases, these competitions suggest anticipated development paths that are consistent with computer movement abilities at level 3 being refined and broadly applied by 2030.

5. Step three: Using current computer ability levels to project future occupational change

The third step in the approach is to apply the assessments of computer ability levels to the current mix of occupational ability requirements to project changes in occupations as these computer abilities are refined and applied over the next couple decades. For this step, the current ability levels from the research discussed in the previous section are used as the projections of broadly applied computer abilities in 2030. Following the pilot results of the last section, this section specifically assumes that computer systems in 2030 will be widely used for language up to level 4 on the O*NET scales, reasoning up to level 5, and vision and movement up to level 3.

The pilot version of the third step involves three simplifying assumptions that transform the occupational projection into a mechanical calculation. First, the pilot approach assumes that any occupation will be automated if it jointly requires language no greater than level 4, reasoning no greater than level 5, and vision and movement no greater than level 3. Second, it assumes that no occupation will be automated if its current ability requirements exceed these limits for one or more of the four groups. Third, it assumes that the relative employment in the occupations that are not automated will stay the same, so that new employment levels can be obtained by simply increasing employment by the same factor across all non-automated occupations. Obviously, a full scale analysis can relax these mechanistic assumptions, but for this pilot effort they allow a plausible and feasible way to quickly estimate the likely scale of the resulting occupational change.

Table 8 shows the results of this simple projection, giving the percentage of current employment that computers would take over in each of the 22 non-military major occupational groups. Across all groups, the projection is for computers to take over occupations representing 60 percent of current employment. The table shows that over two-thirds of employment would be displaced in occupations related to education, food preparation, cleaning, personal care and service, sales, and office and administrative support. At the other end of the spectrum less than one-third of employment would be displaced in occupations related to science and engineering, law, healthcare, protective service, and installation and repair.

Such wholesale displacement of employment would have a large impact on the mix of abilities required in the remaining occupations. Table 9 shows the projected distribution of employment across levels for the four ability groups, which can be compared with the current distribution shown in Table 4. The proportion of the workforce requiring abilities at level 5 or higher rises from 21 to 52 percent for language and from 15 to 36 percent for reasoning. Similarly for vision and movement abilities, the proportion of the workforce requiring abilities at level 4 or higher rises from 17 to 44 percent for vision and from 18 to 45 percent for movement. The comparisons indicate substantial increases in required ability levels for the human workforce.

6. Potential policy issues raised by the prospect of discontinuous change in skill demand

The above pilot analysis shows the feasibility of carrying out a systematic review of current computer research abilities and aligning the conclusions of that review with data on the distribution of ability requirements in the workforce. Obviously, the full analysis has not been implemented, so the results in the previous three sections should be viewed as only suggestive of the possible results of carrying out the full approach. This section briefly discusses some of the policy issues raised by the pilot results as a way of pointing out the value of carrying out the full approach. The potential policy issues involved are serious and are likely to require substantial time to address adequately.

a. Employment and education

Full employment is the immediate policy issue raised by the prospect of substantial displacement of human workers. Of course, there are transition costs for large portions of the workforce to change occupations, but the issues involved in substantial occupational turnover are hardly new. Far more troubling is the prospect that the skills of the workforce will need to be substantially increased. Table 9 suggests that the portion of the workforce capable of performing work at the higher levels of language and reasoning may need to double.

It is an open question whether we can move the full workforce to the projected ability distribution over the next couple decades. It would presumably require unprecedented investments in education for both children and adults to bring the distribution of workforce abilities to the level suggested by Table 9. Such a commitment to education would need to be put in place in the next decade.

b. Competitiveness

Separate from the concern about increasing the skill levels of the workforce in order to maintain employment, there are the investment decisions that government and industry will need to make for the U.S. to stay economically competitive. As available computer abilities increase, the tradeoffs between investing in human expertise and investing in computer expertise will change. For example, the systems summarized in Table 7 include three focused on advanced areas of reasoning in the natural sciences (*cef*). At some point, research-based industries will need to begin to shift their lower level scientific work to computers in order to maintain competitiveness and this will alter the investments currently being made to produce and develop the workforce that currently performs that lower level scientific work.

From a policy point of view, there will be some tension between policies that are pushing computer investment forward for competitiveness and those that are pushing education forward in order to minimize the resulting employment problems. As employment displacement grows it would be reasonable to expect growing resistance to policies that further promote investment in computerization.

c. Economic growth

The large-scale displacement of human workers by computers for a substantial range of skills could cause substantial increases in economic growth. A model of the resulting economic transition suggests that a displacement of the size suggested by the pilot analysis would set the economy on a growth path towards an economy that is an order of magnitude larger than the current economy (Elliott, 1998). Thus the indications we have seen over the past decade of the role computers are playing in economic growth could be early signals of substantial further changes in the decades ahead. If such large levels of growth are plausible, they will require a rethinking about other policy issues—such as social security and global warming—that are influenced by our expectations regarding long-term economic growth.

d. Unemployment and education

In the current economy, it is still the case that most human workers with low skill levels have better language, common sense reasoning, vision, and movement abilities than most current computers do. However, the pilot analysis suggests that this will not be the case in 2030. Certainly workers at lower skill levels will see their wages fall relative to others as computers become capable of providing some of these skills. However, as computer abilities improve, there will come a point when low wages will be insufficient to make a low skill worker competitive with a computer. Today, human workers could not compete with computers for carrying out long division even if they were willing to work for free; we should expect to begin to see that kind of complete displacement in other skill areas as computer abilities improve. So if education is unable to move the ability distribution of the workforce up to the mix suggested by Table 9, there may be many people who will become essentially unemployable over the next couple decades.

However, even if we manage with full employment for the next couple decades, it is important to remember that computer technology is a moving target that will continue to change. Table 9 shows a projection of the results of the broad application of computer abilities by 2030, but a projection showing the broad application of computer abilities by 2040 would be even greater. As long as computer abilities continue to improve, we should expect that the skill requirements for the human workforce will continue to shift up. With this steadily moving target, there will come a time when we are simply unable to move human skills up quickly enough to keep the full workforce employed. At that point, we will need to turn towards policies that address unemployment in a much more fundamental way.

Given the progress of computer abilities over the past 50 years, it is reasonable to expect that computers will surpass human workers in effectively all occupational skill areas during the course of the 21st century. Indeed, there appears to be a significant chance that this could happen before the century is half over (Moravec, 1999). As we approach this point, we will need to refocus our educational, economic, and social policies from their current fixation on employment. Of necessity, this shift will begin with policies that increase the levels of income support for displaced workers who appear to be essentially unemployable. Over time, such support will continue to increase—in length and generosity—as larger portions of the current workforce begin to have trouble competing with computers. Eventually, we will need to sever the link between income and work completely and develop educational and social policies that help people find satisfying and meaningful activities outside the structure of paid employment.

7. Next steps

The pilot effort described in this paper can provide only a suggestion of the results of a full scale effort to understand computer abilities in the current research literature through the lens of the abilities humans use at work. The intent is not to suggest that we currently know how well current computer abilities align with the abilities of the human workforce, but only to suggest that it is possible to know this and that there are a number of important educational, economic, and social policies that could be substantially affected by such knowledge.

The paper notes a number of ways that this pilot effort is a coarse approximation of the full analysis that should be conducted to adequately assess current computer abilities in the research literature, compare them with current occupational requirements, and then project the economic transition as those computer abilities are refined and broadly applied over the next couple decades. An adequate projection will require an effort that marshals a broad range of expertise in computer science, psychology, and the social sciences. Furthermore, it is not enough to simply focus on projecting computer abilities and their impact on the workforce; the potential for a substantial change in both the skill mix of the workforce and the institution of work itself calls for an accompanying research effort related to the educational, economic, and social policies that will be necessary in response.

Over the course of the 21st century, the impact of computers on work will be far-reaching and have profound implications for many of society's most important institutions. It is time to mount a serious and sustained effort to understand and prepare for this change.

References

Aggour, Kareem S., Bonissone, Piero P., Cheetham, William E., and Messmer, Richard P. (2006). "Automating the underwriting of insurance applications." *AI Magazine*, 27(3), 36-50.

Barbuceanu, Mihai, Fox, Mark S., Hong, Lei, Lallement, Yannick, and Zhang, Zhongdong, (2004). "Building agents to serve customers." *AI Magazine*, 25(3), 47-60.

Buchanan, Bruce G., and Livingston, Gary R., (2004). "Toward automated discovery in the biological sciences." *AI Magazine*, 25(1), 69-84.

Chang, Kenneth, (2007). "Researchers go molecular in design of a denser chip." *The New York Times*, January 25.

Elliott, Stuart W., (1998). "Computer technology, human labor, and long-run economic growth." Carnegie Mellon University, Heinz School Working Paper 98-23.

Forbus, Kenneth D., and Hinrichs, Thomas R., (2006). "Companion cognitive systems: A step toward human-level AI." *AI Magazine*, 27(2), 83-95.

Friedland, Noah S., Allen, Paul G., et al., (2004). "Project Halo: Towards a digital Aristotle." *AI Magazine*, 25(4), 29-48.

Gopal, Kreshna, Romo, Tod D., et al. (2006). "TEXTAL: Crystallographic protein model building using AI and pattern recognition." *AI Magazine*, 27(3), 15-24.

Jarvis, Peter A., Lunt, Teresa F., and Myers, Karen L. (2005). "Identifying terrorist activity with AI plan recognition technology." *AI Magazine*, 26(3), 73-81.

Lenat, Douglas B. (1995). "Cyc: A large-scale investment in knowledge infrastructure." *Communications of the ACM*, 38(11).

Lenat, Douglas B. (2006). "Computers versus common sense." Lecture given at Google, May 30, 2006. Available at http://video.google.com/.

Lima, Pedro, Custodio, Luis, et al. (2005). "RoboCup 2004 competitions and symposium: A small kick for robots, a giant score for science." *AI Magazine*, 26(2), 36-61.

Moravec, Hans, (1999). *Robot: Mere machine to transcendent mind*. New York, NY: Oxford University Press.

Nagel, Hans-Hellmut, (2004). "Steps toward a cognitive vision system." *AI Magazine*, 25(2), 31-50.

Peterson, Norman G., Mumford, Michael D., Borman, Walter C., Jeanneret, P. Richard, and Fleishman, Edwin A. (1995). *Development of Prototype Occupational Information Network (O*NET) Content Model*. Utah Department of Workforce Services.

Sebe, Nicu, and Lew, Michael S., (2003). *Robust computer vision: Theory and applications*. Dordrecht: Kluwer Academic Publishers.

Shah, Purvesh, Schneider, David, et al. (2006. "Automated population of Cyc: Extracting information about named entities from the web." Proceedings of the 19th International FLAIRS Conference.

Simmons, Reid, Goldberg, Dani, et al. (2003). "GRACE: An autonomous robot for the AAAI Robot Challenge." *AI Magazine*, 24(2), 51-72.

Swartout, William, Gratch, Jonathan, et al. (2006). "Toward virtual humans." *AI Magazine*, 27(2), 96-108.

Thrun, Sebastian, (2006). "A personal account on the development of Stanley, the robot that won the DARPA Grand Challenge." *AI Magazine*, 27(4), 69-82.

Table 1: O*NET Sets of Descriptors for Workers and Occupations[1]

Worker Characteristics
- Abilities
- Occupational Interests
- Work Values
- Work Styles

Worker Requirements
- Basic Skills
- Cross-Functional Skills
- Knowledge
- Education

Experience Requirements
- Experience and Training
- Basic Skills – Entry Requirement
- Cross-Functional Skills – Entry Requirement
- Licensing

Occupation-Specific Information
- Tasks
- Tools and Technology

Occupational Requirements
- Generalized Work Activities
- Detailed Work Activities
- Organizational Context
- Work Context

---

[1] Source: Description of the O*NET Content Model, http://www.onetcenter.org/content.html, accessed 4/13/07.

Table 2: O*NET Abilities Scales

Cognitive
- Verbal
  - Oral comprehension*
  - Written comprehension*
  - Oral expression*
  - Written expression*
- Idea generation and reasoning
  - Fluency of ideas
  - Originality*
  - Problem sensitivity*
  - Deductive reasoning*
  - Inductive reasoning*
  - Information ordering
  - Category flexibility
- Quantitative
  - Mathematical reasoning*
  - Number facility
- Memory
  - Memorization
- Perceptual
  - Speed of closure*
  - Flexibility of closure*
  - Perceptual speed*
- Spatial
  - Spatial orientation*
  - Visualization*
- Attentiveness
  - Selective attention
  - Time sharing

Psychomotor
- Fine manipulative
  - Arm-hand steadiness
  - Manual dexterity*
  - Finger dexterity*
- Control movement
  - Control precision*
  - Multilimb coordination*
  - Response orientation
  - Rate control*
- Reaction time and speed
  - Reaction time
  - Wrist-finger speed
  - Speed of limb movement

Physical
- Physical strength
  - Static strength
  - Explosive strength
  - Dynamic strength
  - Trunk strength
- Endurance
  - Stamina
- Flexibility, balance, coordination
  - Extent flexibility
  - Dynamic flexibility
  - Gross body coordination*
  - Gross body equilibrium*

Sensory
- Visual
  - Near vision
  - Far vision
  - Visual color discrimination
  - Night vision
  - Peripheral vision
  - Depth perception*
  - Glare sensitivity
- Auditory and speech
  - Hearing sensitivity
  - Auditory attention
  - Sound localization
  - Speech recognition
  - Speech clarity

* Used in pilot study

Table 3: O*NET Abilities Scales Used in Pilot Study with Questionnaire Anchoring Tasks[1]

Table 3a: Language Abilities

Scale	Low (Level = 2)[2]	Medium (Level = 4)	High (Level = 6)
Oral comprehension	Understand a television commercial	Understand a coach's oral instructions for a sport	Understand a lecture on advanced physics
Written comprehension	Understand signs on the highway	Understand an apartment lease	Understand an instruction book on repairing missile guidance systems
Oral expression	Cancel newspaper delivery by phone	Give instructions to a lost motorist	Explain advanced principles of genetics to college freshmen
Written expression	Write a note to remind someone to take food out of the freezer (Level = 1)	Write a job recommendation for a subordinate	Write an advanced economics textbook

Table 3b: Reasoning Abilities

Scale	Low (Level = 2)	Medium (Level = 4)	High (Level = 6)
Originality	Use a credit card to open a locked door	Redesign job tasks to be interesting for employees	Invent a new type of man-made fiber
Problem sensitivity[3]	Recognize that an unplugged lamp won't work	Recognize from the mood of prisoners that a prison riot is likely to occur	Recognize an illness at an early stage of a disease when there are only a few symptoms
Deductive reasoning	Know that a stalled car can coast downhill	Decide what factors to consider in selecting stocks (Level = 5)	Design an aircraft wing using principles of aerodynamics
Inductive reasoning	Decide what to wear based on the weather report	Determine the prime suspect based on crime scene evidence	Diagnose a disease using results of many different lab tests
Mathematical reasoning	Determine how much 10 oranges will cost when they are priced at 2 for 20 cents (Level = 1)	Decide how to calculate profits to determine the amounts of yearly bonuses	Determine the mathematics required to simulate a space craft landing on the moon
Visualization	Imagine how to put paper in a typewriter so that the letterhead comes out on top	Follow a diagram to assemble a metal storage cabinet	Anticipate opponent's as well as your own future moves in a chess game

---

[1] Source: O*NET Abilities Questionnaire, http://www.onetcenter.org/questionnaires.html, accessed 4/14/07.

[2] All anchoring tasks are at the typical level—2 for low, 4 for medium, 6 for high—except where noted.

[3] "The ability to tell when something is wrong or is likely to go wrong. It does not involve solving the problem, only recognizing that there is a problem."

Table 3c: Vision Abilities

Scale	Low (Level = 2)	Medium (Level = 4)	High (Level = 6)
Speed of closure[4]	Recognize a song after hearing only the first few notes (Level = 3)	Make sense out of strange handwriting	Interpret patterns on weather radar to decide if the weather is changing
Flexibility of closure[5]	Tune in a radio in a noisy truck	Look for a golf ball in the rough	Identify camouflaged tanks from a high-speed airplane
Perceptual speed	Sort mail according to ZIP codes with no time pressure	Read five temperature gauges in 10 seconds to make sure each temperature is within safe limits	Inspect electrical parts for defects as they flow by on a fast-moving assembly line
Depth perception	Merge a car into traffic on a city street	Operate a crane to move materials from a truck bed to the ground.	Throw a long pass to a closely guarded teammate

Table 3d: Movement Abilities

Scale	Low (Level = 2)	Medium (Level = 4)	High (Level = 6)
Spatial orientation	Use the floor plan to locate a store in a mall	Find your way through a dark room without hitting anything (Level = 3)	Navigate an ocean voyage using only the positions of the sun and stars
Manual dexterity	Screw a light bulb into a light socket (Level = 1)	Pack oranges in crates as quickly as possible	Perform open heart surgery with surgical instruments (Level = 7)
Finger Dexterity	Put coins in a parking meter	Attach small knobs to stereo equipment on an assembly line	Put together the inner workings of a small wrist watch
Control precision	Adjust a room light with a dimmer switch	Adjust farm tractor controls	Drill a tooth
Multilimb coordination	Row a boat	Operate a forklift truck in a warehouse	Play the drum set in a jazz band
Rate control	Ride a bicycle alongside a jogger (Level = 1)	Keep up with a car that changes speed	Shoot a duck in flight (Level = 5)
Gross body coordination	Get in and out of a truck	Swim the length of a pool	Perform a ballet dance
Gross body equilibrium	Stand on a ladder	Walk on ice across a pond	Walk on narrow beams in high-rise construction

---

[4] "The ability to quickly make sense of, combine, and organize information into meaningful patterns."

[5] "The ability to identify or detect a known pattern (a figure, object, word, or sound) that is hidden in other distracting material."

Table 4: Distribution of Employment across Ability Groups

Ability Level	Language Group	Reasoning Group	Vision Group	Movement Group
1	0%	0%	5%	7%
2	2%	2%	26%	31%
3	19%	42%	52%	43%
4	58%	41%	16%	17%
5	20%	14%	1%	1%
6	1%	1%	0%	0%
7	0%	0%	0%	0%

Table 5: Average Ability Ratings for Minor Occupational Groups

Minor SOC Occupational Group		Emp	L	R	V	M
11-1000	Top executives	1.6%	5	5	4	2
11-2000	Advertising, marketing, promotions, public relations, and sales managers	0.5%	4	4	3	1
11-3000	Operations specialties managers	1.1%	5	4	3	2
11-9000	Other management occupations	2.6%	4	4	3	2
13-1000	Business operations specialists	1.7%	5	4	3	2
13-2000	Financial specialists	1.7%	5	4	3	2
15-1000	Computer specialists	2.1%	5	4	3	3
15-2000	Mathematical science occupations	0.1%	5	5	3	2
17-1000	Architects, surveyors, and cartographers	0.2%	5	5	3	3
17-2000	Engineers	0.9%	5	5	4	2
17-3000	Drafters, engineering, and mapping technicians	0.5%	5	4	3	3
19-1000	Life scientists	0.2%	5	5	4	3
19-2000	Physical scientists	0.2%	6	5	4	2
19-3000	Social scientists and related occupations	0.3%	5	5	3	1
19-4000	Life, physical, and social science technicians	0.2%	5	4	4	3
21-1000	Counselors, social workers, and other community and social service specialists	1.1%	4	4	3	2
21-2000	Religious workers	0.4%	5	4	3	2
23-1000	Lawyers, judges, and related workers	0.6%	5	5	3	2
23-2000	Legal support workers	0.3%	5	4	4	2
25-1000	Postsecondary teachers	1.2%	5	4	2	1
25-2000	Primary, secondary, and special education teachers	3.1%	4	4	2	2
25-3000	Other teachers and instructors	0.3%	4	3	2	2
25-4000	Librarians, curators, and archivists	0.2%	4	4	4	3
25-9000	Other education, training, and library occupations	1.0%	4	3	2	2
27-1000	Art and design occupations	0.5%	4	4	3	3
27-2000	Entertainers and performers, sports and related occupations	0.5%	4	4	3	3
27-3000	Media and communication occupations	0.5%	5	4	2	1
27-4000	Media and communication equipment occupations	0.2%	4	4	3	3
29-1000	Health diagnosing and treating practitioners	3.0%	5	5	4	3
29-2000	Health technologists and technicians	1.7%	4	4	4	3
29-9000	Other healthcare practitioners and technical occupations	0.0%	4	4	3	3
31-1000	Nursing, psychiatric, and home health aides	1.5%	4	4	3	3
31-2000	Occupational and physical therapist assistants and aides	0.1%	4	3	2	3
31-9000	Other healthcare support occupations	0.7%	4	3	2	3
33-1000	First-line supervisors/managers, protective service workers	0.1%	4	4	4	3
33-2000	Fire fighting and prevention workers	0.2%	4	4	4	4
33-3000	Law enforcement workers	0.9%	4	4	4	3
33-9000	Other protective service workers	0.9%	4	4	4	3
35-1000	Supervisors, food preparation and serving workers	0.7%	4	4	2	2
35-2000	Cooks and food preparation workers	2.2%	4	3	3	3
35-3000	Food and beverage serving workers	4.0%	3	3	1	2
35-9000	Other food preparation and serving related workers	0.9%	3	2	2	3
37-1000	Supervisors, building and grounds cleaning and maintenance workers	0.3%	4	3	2	2
37-2000	Building cleaning and pest control workers	2.8%	3	3	2	3
37-3000	Grounds maintenance workers	0.9%	3	3	3	3
39-1000	Supervisors, personal care and service workers	0.2%	4	3	3	2
39-2000	Animal care and service workers	0.1%	4	4	3	3

Table 5: Average Ability Ratings for Minor Occupational Groups (cont.)

Minor SOC Occupational Group		Emp	L	R	V	M
39-3000	Entertainment attendants and related workers	0.4%	4	3	2	2
39-4000	Funeral service workers	0.0%	4	3	2	3
39-5000	Personal appearance workers	0.6%	3	3	3	3
39-6000	Transportation, tourism, and lodging attendants	0.2%	4	3	2	3
39-9000	Other personal care and service workers	1.8%	4	4	3	3
41-1000	Supervisors, sales workers	1.6%	4	4	3	2
41-2000	Retail sales workers	6.1%	4	3	3	2
41-3000	Sales representatives, services	0.7%	4	4	3	2
41-4000	Sales representatives, wholesale and manufacturing	1.3%	4	4	3	2
41-9000	Other sales and related workers	0.9%	4	3	2	2
43-1000	Supervisors, office and administrative support workers	1.1%	5	4	3	2
43-2000	Communications equipment operators	0.2%	3	2	2	2
43-3000	Financial clerks	2.8%	4	3	3	2
43-4000	Information and record clerks	3.8%	4	3	3	3
43-5000	Material recording, scheduling, dispatching, and distributing occupations	2.8%	4	3	2	3
43-6000	Secretaries and administrative assistants	3.0%	4	3	2	1
43-9000	Other office and administrative support workers	3.2%	4	3	3	3
45-1000	Supervisors, farming, fishing, and forestry workers	0.0%	4	4	3	3
45-2000	Agricultural workers	0.6%	3	3	3	3
45-3000	Fishing and hunting workers	0.0%	3	2	3	4
45-4000	Forest, conservation, and logging workers	0.1%	2	3	3	4
47-1000	Supervisors, construction and extraction workers	0.5%	4	4	3	3
47-2000	Construction trades and related workers	4.3%	3	3	3	4
47-3000	Helpers, construction trades	0.3%	3	3	2	3
47-4000	Other construction and related workers	0.3%	3	3	3	3
47-5000	Extraction workers	0.1%	3	3	3	4
49-1000	Supervisors of installation, maintenance, and repair workers	0.3%	5	4	4	3
49-2000	Electrical and electronic equipment mechanics, installers, and repairers	0.5%	4	4	3	4
49-3000	Vehicle and mobile equipment mechanics, installers, and repairers	1.3%	4	4	3	4
49-9000	Other installation, maintenance, and repair occupations	1.9%	4	4	3	4
51-1000	Supervisors, production workers	0.5%	4	4	3	3
51-2000	Assemblers and fabricators	1.2%	4	3	3	3
51-3000	Food processing occupations	0.5%	3	3	3	3
51-4000	Metal workers and plastic workers	1.5%	3	4	3	4
51-5000	Printing occupations	0.3%	3	3	2	3
51-6000	Textile, apparel, and furnishings occupations	0.7%	3	3	2	3
51-7000	Woodworkers	0.2%	3	4	3	4
51-8000	Plant and system operators	0.2%	3	4	3	3
51-9000	Other production occupations	1.9%	3	3	3	3
53-1000	Supervisors, transportation and material moving workers	0.3%	4	4	3	3
53-2000	Air transportation occupations	0.1%	4	5	4	4
53-3000	Motor vehicle operators	3.0%	4	4	3	4
53-4000	Rail transportation occupations	0.1%	4	4	3	3
53-5000	Water transportation occupations	0.1%	3	3	4	4
53-6000	Other transportation workers	0.2%	4	3	3	3
53-7000	Material moving occupations	3.5%	3	3	3	3

Table 6: Distribution of Current Employment across Joint Language-Reasoning (LR) and Vision-Movement (VM) Ability Groups

	VM=1	VM=2	VM=3	VM=4	VM=5	VM=6	VM=7
LR=1	0%	0%	0%	0%	0%	0%	0%
LR=2	0%	0%	0%	1%	0%	0%	0%
LR=3	0%	2%	11%	5%	0%	0%	0%
LR=4	0%	13%	34%	11%	0%	0%	0%
LR=5	0%	2%	9%	10%	1%	0%	0%
LR=6	0%	0%	0%	1%	0%	0%	0%
LR=7	0%	0%	0%	0%	0%	0%	0%

Table 7: Summaries of 12 Computer Systems from *AI Magazine*, 2003-2006

a. Aggour et al. (2006), “Automating the underwriting of insurance applications”

A system has been developed and deployed at a large insurance provider to automate the underwriting of long-term care insurance. The system uses input from paper applications that has been summarized. The system can make decisions about cases with no medical issues, reason about medical impairments that cover two-thirds of the applicants with impairments, order additional information—such as a physician summary or face-to-face interview—when required, identify whether the free text entered in the summarized application is critical to the underwriting decision, and refer the application to a human underwriter if a clear decision isn’t indicated. The system initially deals with all summarized applications and renders final decisions for 19%. For the remaining applications that are forwarded to a human underwriter, the system identifies the conflicting recommendations that prevented it from being able to make a decision.

b. Barbuceanu et al. (2004), “Building agents to serve customers”

A system is being developed to provide customer service assistance that combines language use and the ability to reason about business issues. The system combines two types of language processing – an in-depth analytic approach for primary topics and a more approximate similarity approach for secondary topics – in order to deal with a wide range of situations. The system is adapted to handle web, email, and telephone interactions, with appropriate styles of communication for each mode. The system has been evaluated in seven different business applications and correctly addresses customers’ issues roughly two-thirds of the time. The system has been used with a range of customer service topics, including questions about order status, product quality, billing and payment, business policy, and diagnosis of service problems.

c. Buchanan and Livingston (2004), “Toward automated discovery in the biological sciences”

A system is being developed to discover potentially interesting scientific hypotheses in empirical data. This work goes beyond data-mining techniques by using both general heuristics and a number of types of domain knowledge to identify hypotheses. The system has been tested by looking for new hypotheses about conditions for growing crystals for determining the three-dimensional structure of proteins, using a dataset of 2,225 cases of experimental results growing crystals. Of the 575 hypotheses generated by the system, 92 (16%) were judged to be of sufficient novelty and significance to justify further investigation. The generality of the approach was evaluated by applying the system to 930 cases of rehabilitation after a medical disability, which resulted in the discovery of 299 hypotheses, 26 (9%) of which were judged to be of sufficient novelty and significance to justify further investigation.

Table 7: Summaries of 12 Computer Systems from *AI Magazine*, 2003-2006 (cont.)

d. Forbus and Hinrichs (2006), "Companion cognitive systems: A step toward human-level AI"

A system is being developed that can learn about a domain and interact with humans to provide useful suggestions. Mirrored from findings about human psychology, the system reasons and learns by analogy. Input to the system is done with sketches and concept maps. One application of the system has been in the area of everyday physical reasoning, such as a problem to determine which of two wheelbarrows will be easiest to lift drawn from a test of mechanical reasoning ability. Another application of the system has been in the area of military tactical decision games, drawn from *USMC Gazette*, which involve a map showing terrain and force layouts and a written description of the military situation. A third application of the system involves a civilization game, which is similar to the military game except that it involves economic and political decisions in addition to military decisions. In all these applications, the system provides potentially relevant advice by retrieving analogous cases that could be used to solve the problem. The system is being evaluated by its ability to recommend useful analogies, but the article doesn't report any comparisons to human performance on the problems.

e. Friedland et al. (2004), "Project Halo: Towards a digital Aristotle"

A challenge was conducted with three groups working in the area of knowledge representation and reasoning. The task involved building a system with the knowledge contained in 70 pages of an advanced placement (AP) chemistry textbook and the ability to use that knowledge to answer novel questions and generate explanations in English. Each team had four months and $700,000 to develop their system. The resulting systems were evaluated using an AP-style test involving multiple choice, fill-in-the-blank, and short essay questions that were graded by three chemistry professors. For all types of questions, the systems provided both an answer and a justification of that answer. Given the short timeframe for the project, the teams were allowed to encode the questions in their own system's formal knowledge representation language, as long as the formal representation included no extra information. All three systems achieved average correctness scores above 40%, comparable to the mean score for students taking the AP chemistry exam. The challenge is a pilot effort for a project that would encode a broad range of scientific knowledge in a system for educational instruction and interdisciplinary research assistance.

Table 7: Summaries of 12 Computer Systems from *AI Magazine*, 2003-2006 (cont.)

f. Gopal et al (2006), "TEXTAL crystallographic protein model building using AI and pattern recognition"

To determine the structure of proteins, x-ray crystallography is used to generate electron density maps that are then fit to possible underlying atomic structures. The process of creating a refined atomic model is complex and can take weeks or months for an expert crystallographer. It can take several days for a crystallographer to fit a potential model to each of the electron density maps used. A computer system has been developed for building atomic models that is fully automated and can work with electron density maps of average quality. The system mimics expert approaches that look for distinctive chains forming the backbone of the protein and then identify densities of side chains. The system combines a number of AI heuristic search and knowledge-based techniques. It produces a reasonable model in a couple hours that can then be refined by an expert. The system typically identifies roughly 80-90 percent of the backbone correctly and 50 percent of the side chains. The system is being used in over 100 crystallography labs and is being regularly improved.

g. Jarvis et al. (2005), "Identifying terrorist activity with AI plan-recognition technology"

A system is being developed to recognize patterns in a set of potentially suspicious facts that could indicate a terrorist plan. An example of such a potentially suspicious fact would be information about a potentially suspicious person obtaining information about a potential terrorist target such as a dam. The system uses AI plan-recognition techniques and a set of templates of types of attacks to combine isolated facts into coherent evidence of an actual plan. The input to the system will be the output of data-mining software, though development has so far used synthetic data. The computing time to recognize a plan increases exponentially with as the number of potentially surprising facts increases – as a result, current PC processing power and algorithms effectively limit the use of the system to a maximum of 20 facts. The article doesn't provide a comparison for people's ability to recognize potential terrorist plans in such sets of potentially suspicious facts.

h. Lima et al. (2005), "RoboCup 2004 competitions and symposium"

The RoboCup Federation hosts an annual competition and symposium, with the objective of fielding a team of humanoid robot soccer players that by 2050 will win a game against the most recent world cup champions. RoboCup 2004 included over 700 robots in 346 teams, participating in nine different leagues, with robots of different sizes and configurations that present different challenges. The humanoid league was started in 2002, with tests of skills involving walking, kicking a ball into a goal, diving by a goalkeeper, and passing a ball between two players. Movements such as walking have improved but are still significantly slower than for humans, while visual recognition isn't robust. Rules for RoboCup 2006 specify a series of 10-minute matches between two teams, each involving one or two humanoid robots (RoboCup Federation 2006).

Table 7: Summaries of 12 Computer Systems from *AI Magazine*, 2003-2006 (cont.)

i. Nagel (2004), “Steps toward a cognitive vision system”

A computer vision system has been developed over 30 years with the goal of verbally describing the activities that are taking place in a visual scene. The system has been focused on describing the movement of vehicles on a street, with the development of robust approaches for identifying vehicles in successive visual images based on models of various types of vehicles. The system also identifies street lanes, but does not identify other objects in a street scene, such as pedestrians. The system is able to identify vehicles as they are moving behind occlusions. Typical vehicle movements—such as approaching or leaving a location—have been defined so that they can be identified from the visual images. These descriptions of vehicle movements are then transformed into natural language.

j. Simmons et al. (2003), “GRACE: An autonomous robot for the AAAI Robot Challenge”

The AAAI Robot Challenge task is for a robot to attend the AAAI national conference as a participant, including finding the registration booth, registering for the conference, interacting with people at the conference for information and personal exchanges, and giving a technical talk and answering questions. The task requires moving around a complex environment and interacting with people through both language and vision. The GRACE system was entered in 2002 and included four computers running systems for speech recognition, language understanding, stereo vision, color vision processing, facial animation (on a screen), navigation and localization. Numerous problems were encountered during the actual events, as systems such as speech and gesture recognition that had worked relatively reliably in the lab had difficulty with the confusion of the conference setting. In the elevator-riding task, GRACE required that humans hold the doors open because it didn’t move quickly enough. In addition, it had no way of determining what floor it was on. The system was able to successfully locate specially designed registration booths, but it then pushed its way into the middle of the waiting line. The system also became disoriented when moving down a hall after tables had been placed in the hallway and a crowd of onlookers prevented it from being able to locate the walls. The system gave its talk using PowerPoint slides and a conceptual map of the topics to be covered; it used the bulleted points on the slides to locate relevant points in its conceptual map and then for each bullet explained the most important and relevant point that had not yet been explained, allowing it to handle questions that arose during the talk without repeating material unnecessarily.

Table 7: Summaries of 12 Computer Systems from *AI Magazine*, 2003-2006 (cont.)

k. Swartout et al. (2006), “Toward virtual humans”

A system is being developed to provide interactive virtual humans for a virtual reality training system. The virtual humans include modules that understand and generate speech and that reason about tasks, social situations and emotions. The training scenario that has been implemented involves a lieutenant (the student) who must decide what to do when he unexpectedly comes across an accident between a military and a civilian vehicle while driving to a meeting. Reasoning about emotion plays important roles in the system, for such tasks as choosing emotionally appropriate language and understanding imprecise language. The system includes a vocabulary of a few hundred words for speech recognition and a vocabulary of a thousand words for speech generation. The virtual humans include models describing about 40 tasks and 150 properties of the world. In interactions with human students, the system responds to an utterance by the student in about three seconds, including understanding the utterance, updating the system’s model of the dialogue and the student’s emotional state, choosing how to respond, and producing the appropriate voice output and accompanying gestures.

l. Thrun (2006), “A personal account of the development of Stanley, the robot that won the DARPA Grand Challenge”

DARPA announced a Grand Challenge in 2004 that called for robotic vehicles to drive independently along a 132-mile-long course in the Mojave Desert in less than 10 hours. Five vehicles successfully completed the challenge in 2005, with the winning system completing the course in less than seven hours driving a VW Touareg SUV. The winning team assembled their system in a year’s time with a team of Stanford professors and students. The vehicle uses a laser mapping system to detect road surface and obstacles near the vehicle combined with a computer vision system for detecting the continuation of the drivable surface farther from the vehicle. This combination allows the vehicle to adapt to different road types and to drive up to 38 mph with the longer-range computer vision system while slowing to 25 mph when the computer vision indicates an upcoming change in road surface that requires the shorter-range laser mapping system to negotiate (Montemerlo et al. 2006). The vehicle uses learning algorithms and probabilistic reasoning to adapt to changing conditions and errors in its own sensors. Although the vehicle successfully avoids obstacles and executes turns, it sometimes makes these changes later than a human driver would in the same situation, resulting in more abrupt corrections. The Grand Challenge didn’t require driving in the presence of other vehicles or reacting appropriately to traffic signs and regulations—these additional challenges will be added in the DARPA Urban Challenge, which will be held in late 2007.

Table 8: Projected Displacement by 2030 in SOC Major Occupational Groups

Major Occupational Group		Percent of Total Employment	Percent Displaced Within Group
11-0000	Management	6%	41%
13-0000	Business and financial operations	3%	32%
15-0000	Computer and mathematical science	2%	21%
17-0000	Architecture and engineering	2%	11%
19-0000	Life, physical, and social science	1%	10%
21-0000	Community and social services	2%	36%
23-0000	Legal	1%	6%
25-0000	Education, training, and library	6%	74%
27-0000	Arts, design, entertainment, sports, and media	2%	50%
29-0000	Healthcare practitioners and technical	5%	10%
31-0000	Healthcare support	2%	29%
33-0000	Protective service	2%	16%
35-0000	Food preparation and serving related	8%	88%
37-0000	Building and grounds cleaning and maintenance	4%	78%
39-0000	Personal care and service	3%	81%
41-0000	Sales and related	11%	93%
43-0000	Office and administrative support	17%	90%
45-0000	Farming, fishing, and forestry	1%	43%
47-0000	Construction and extraction	5%	39%
49-0000	Installation, maintenance, and repair	4%	12%
51-0000	Production	7%	53%
53-0000	Transportation and material moving	7%	64%
Total		100%	60%

Table 9: Projected Distribution of Employment across Ability Levels in 2030

Ability Level	Language Group	Reasoning Group	Vision Group	Movement Group
1	0%	0%	0%	7%
2	3%	2%	10%	22%
3	14%	17%	46%	25%
4	30%	45%	42%	43%
5	50%	34%	2%	2%
6	2%	2%	0%	0%
7	0%	0%	0%	0%