The robot in the Garden

Nikolaus Correll
Department of Computer Science, University of Colorado at Boulder

In: Michael G. Lee, Kenneth L. Helphand (Ed.): Technology and the Garden, Harvard University Press, 2014.

Since the first deliberate sowing and harvesting around ten thousand years ago, agricultural techniques have been industrially streamlined and transformed the vegetable garden from a necessary part of every household to a luxury hobby in the developed world. Efficiency has increased exponentially with the industrialization of the agricultural sector in the 20iest century and lead to a system that has become a major challenge to our eco-system, economy, health-care system, and national security[i]. At the same time the gap between the garden and the machine postulated by Leo Marx[ii] has been ever increasing, disconnecting food production from consumption. The working thesis of this paper is that these challenges are driving technological evolution toward small-scale autonomous systems that provide care to individual plants. These robotic systems might remedy not only the challenges of the current agricultural system, but also bring agriculture closer to humans: Being able to tend to individual plants on a need-basis will enable decentralized production of food in ecologically sustainable poly-cultures, optimal assignment of nutrients to each plant, and a reduction of overall pesticide consumption. This transition comes with labor requirements that cannot be provided by current automated systems and the current level of technology, and therefore require additional research in perception, manipulation, control and plant sciences.

This paper is organized as follows: After a brief history on agriculture, we will describe the challenges of the current agricultural system at the time of writing. We will then review the state-of-the art in perception, manipulation, and artificial intelligence techniques with respect to autonomous agriculture, and outline the requirements for an autonomous precision agriculture system that can tend to poly-cultures. The paper is concluded with a discussion on technical and societal challenges that such a system would entail.

Introduction

Since the advent of agriculture around 12,000 years ago, humans have developed a highly sophisticated system for global food production with the most rapid technological advances occurring during the second half of last century. In the 1920ies, agriculture has not only adopted new machinery, but also the financial, cultural, and ideological apparatus of industrialism[iii]. This process has led to ever increasing farm equipment, modern plant breeding programs, the use of synthetic fertilizers, delivery of water via irrigation systems, and the use of pesticides to control crop herbivory, which have all contributed to a tremendous increases in crop yield. For example, corn yields in the US have increased approximately 400%-500% from 1940 to 1997[iv]. Together with consolidation of small farms into larger ones, the number of people that a single farmer in the US provides for increased 9 fold from 15.5 people per farm in 1950 to 140 people per farm in 1997[v]. These developments came at tremendous environmental costs. Increased fertilization has led to excess nitrogen and phosphorus in the water systems impacting both human health and the integrity of aquatic ecosystem and increases in nitrous oxide production (a potent greenhouse gas)[vi]. More recently, there is also an increasing awareness that the global supply of phosphorus, which is a non-renewable resource but an essential plant nutrient contained in most synthetic fertilizers, is expected to peak mid-century and decline thereafter[vii]. Also, cultivation of just a few crops (corn, soybeans, hay and wheat make up 68% of farm land in the US[viii]) in ever increasing mono-cultures and the resulting lack of plant diversity make these systems vulnerable to large scale pest outbreaks. Finally, the high specialization on certain crops in different parts of the US - which is the result of industrial streamlining the process - requires considerable transportation cost, and might become infeasible with increasing cost of oil. Taken together, there is a critical need to develop agricultural practices that deliver water and nutrients in a manner that minimizes losses from these systems, while creating a an ecologically resilient agricultural system that can withstand or quickly recover from disturbances such as pest outbreaks.

Figure 1: Artist impression of an indoor precision agriculture system that is co-located with people.

This paper argues that advances in robotics can decrease detrimental effects of farming by precise administration of water and nutrients and inter-cropping, while bringing agriculture closer to consumers, and outlines challenges that lie on the way to implement this vision. Besides being able to provide individual plants with the required resources on a need-basis, small-scale robotic platforms are not limited to operate on fields, but could also re-claim urban environments that are currently deprived from agriculture and horticulture such as within offices, shopping malls or roofs. A concept drawing of a team of robots cultivating plants in a shopping mall is shown in Figure 1.

State-of-the-Art

The vision of plant-level precision agriculture by autonomous agents is enabled by four complementary factors: First, the increasing cost of non-renewable resources such as oil and phosphorous; second, the current trend in agriculture to precision agriculture[ix]; third, the ability of so-called "companion plants"[x] to benefit each other when planted in close proximity, and fourth recent advances in computing, perception, actuation, and control that have enabled autonomous systems with a high level of robustness[xi] and manipulation dexterity[xii]. The remainder of this section focuses on particular instances of sensing, actuation, computation, and systems thereof, which are specific to the agricultural domain.

Sensing

Sensing is the process of recording data from the environment with the goal to infer the status of a system under observation. In an agricultural application, one of the prime observables are water, light and a plant's developmental stage. Given knowledge on how the plant "works", these observations can be used to infer what the plant needs. In addition to soil moisture and ambient light sensors, other sensing modalities include vision, depth images, infrared spectroscopy, and touch, to name a few.

Ultimately, sensor data support a decision making process. For example, local soil moisture measurements can help to deploy water in a resource-optimal way[xiii], whereas recordings of the reflected infrared spectrum might reveal plant stress and serves as an indicator for nutrient requirements[xiv]. A consumer example that focusses exclusively on automated sensing is the product "EasyBloom", a pen-shaped, water-proof device with a moisture, light and temperature sensor[xv]. This data is then used to match the actual humidity, light, and temperature requirements of a variety of plants from a database to help the gardener to choose plants to grow in this specific location.

One of the richest source of information is vision. General solutions for extracting information such as the maturity of fruits, the location of stems and leaves, or their health that are robust to environmental and lighting conditions, for extracting information remains a hard problem[xvi], however. Nevertheless, solutions for specific plants and environments have demonstrated their potential for being used in production environments, for example the distance between branch junctions of the cotton plant can serve as cue for its irrigation schedule[xvii] , the size of the shaded area under a grapevine canopy is related to its water use[xviii], or thermal imaging can be used to assess apple maturity in an orchard. Recently, a new class of sensors has emerged that can provide imagery and depth data at high resolution and speed. Driven mainly by the gaming industry (with the Microsoft Xbox Kinect the most prominent example), it is now possible to quickly perceive the full 3D geometry of a plant's canopy at low cost.

Actuation

The number of distinct actions that are required from planting to harvesting are enormous and versatile, classically requiring distinct equipment with specific effectors. Advances in sensing (see above) have led to autonomous solutions for *harvesting* specialty crops such as apples[xix], cherries[xx], cucumbers[xxi], tomatoes[xxii], melons[xxiii], mushrooms[xxiv], and strawberries[xxv], among others. Another important task is *trimming*, that is clipping of leaves and branches so that the plant grows in a desired shape. For example, stems can be located in 3D using stereo vision, which can then guide a clipping tool for trimming roses[xxvi], or to prune grape vines, which is commercialized by Vision Robotics[xxvii]. Also automatic *weeding* is an active area of research[xxviii]. Little research has been conducted on other important manipulation functions, such as tying and untying plants from guiding poles or excavating entire plants.

While economic drivers seem yet not to be sufficient for broad deployment of autonomous systems for tending to specialty crops, research into enabling technologies such as low-cost robotic arms with high dexterity, is fueled by applications not related to agriculture, such as tele-operation in space[xii] or robotic surgery[xxix].

Computation

The availability of powerful, cheap, and light-weight computation for agricultural purposes is a result of Moore's law who predicted the density of integrated circuits to double every 2 years. This law has persisted till today and is mainly driven by the proliferation of computing in society and the resulting strong economical drivers.

In addition to the computation, signal processing, and control involved in sensing and actuation, gardening itself is a computational problem. At the plant level, an individual plant can be understand as a control system that reacts to *input* of nutrients, water, light and trimming, by vegetative and fruit growth. For example, tomato plants exhibit mostly vegetative, i.e. green mass, growth when only provided with nitrogen, but require potassium and phosphorus for fruit growth[xxx]. Decisions on what kind of input a plant needs based on observations acquired using sensors require an in-depth expert knowledge of plant dynamics, which can be obtained by systematic studies that record growth parameters vs. nutrient/water/air intake[xxxi] and are dominated by Liebig's principle of the minimum[xxxii], i.e. plant growth is dominated by the most limited resource. Gathering systematic data on how different control parameters affect plant growth allows deriving precise models of plant dynamics. For example, the EPIC crop growth model[xxxiii] provides differential equation-based models for a series of crops (wheat, corn, rice, sunflowers, soybeans, barley) and models the conversion to biomass, the division of biomass into roots, above ground mass, economic yield, and root growth as a function of water use, nutrient uptake and solar radiation. This model requires 22 parameters that define each species and it is governed by 65 equations. Although this complexity lets one assume that the model is comprehensive, it is still far from modeling the complexity of a real living system that is affected by locally varying environmental conditions, and therefore still requires feedback when used in a control context. Being able to get a thorough understanding how a plant will react when presented with varying resources under different environmental conditions and given its specific history is crucial for optimally controlling a plant, however. We believe that progress in this direction, which is currently mainly led by researchers in agriculture and plant biology, will be enabled more and more by robotic technology, for example by automating measurements by physically moving a probe or by relying on recent progresses in sensing.

Systems

Together sensing, actuation, and computation potentially allow integrated systems that respond to actual plant needs with appropriate actuation, much in the way a human gardener could do. Systems that reliably do this are not existent at the time of writing or focus only on very specific aspects of the process in specialized, large-scale agricultural settings involving a single plant species. In order to exemplify the state of the art, we will describe two small-scale robotic gardening systems: the Telegarden[xxxiv] that was operational from 1995 to 2004 and allowed people from all over the world to plant seeds and tend to the plants via the Internet, and the distributed robotic garden project[xxxv] that was started in 2008 and aims at fully autonomous growth of tomato plants, that is watering the plants, assessing plant status, and picking the fruits.

The TeleGarden (USC)

The Telegarden, has been developed by researchers at the University of Southern California (Prof. George Bekey and Prof. Ken Goldberg), together with performance artists (Joseph Santarromana). It is primarily an art installation that allows internet users to view and interact with a remote garden filled with living plants. Members can plant, water, and monitor the progress of seedlings via the tender movements of an industrial robot arm. The Telegarden went online in summer 1995 and attracted over 9000 visitors to help cultivating it in the first year. The garden was then moved to the lobby of the Ars Electronica Center in Austria, where it remained till 2004. The system consists of an industrial robot arm that is statically mounted in the center of a circular pot filled with soil, see also Figure 2. The robot arm's end-effector is equipped with a camera, a watering hose, and a mechanism to deposit seeds into the soil, which are networked with the internet. Users have the ability to login to the system from anywhere in the world, control the position of the robot arm, observe the field through the camera, and trigger watering and seed deposition.

Figure 2: The Telegarden (1995-2004, networked art installation at Ars Electronica Museum, Austria.) Co-directors: Ken Goldberg and Joseph Santarromana Project team: George Bekey, Steven Gentner, Rosemary Morris Carl Sutter, Jeff Wiegley, Erich Berger. (Photo by Robert Wedemeyer).

The system thus provides only two distinct forms of *actuation*, distributing water and depositing seeds. It has a single mode of *sensing*, namely vision. All decisions being made by online visitors based on their perception are limited to these two modes of actuation and are exclusively informed by visual feedback through the robot's sensor. There are three important points with respect to further automation of this system: First, it is irrelevant for the operation of the garden, if these decisions are made by people – who can download pictures of the garden to their computer, look at them, and send back actuation commands – or by a computer that is programmed with some kind of agenda when and where to plant plans, and that knows when which plant needs water and how much. Second, making all these decisions, that is understanding what one sees on the images and planning the growth cycle is still – in 2012 – a hard technical problem. Third, even though the

Telegarden is fully remote controlled, the robotic system performs only a fraction of the tasks that a production system really requires to solve: external, human supervision is needed to remove old or dead plants and weeds, or deploy pesticides in case of a pest infection, among others.

Being an artistic installation, the Telegarden also shed light on how technology can help to bring people closer to people and bridge the gap between the machine and the garden. Ken Goldberg describes the project as follows: "*In the past 10 years since the garden went online we've had over 10,000 members that have registered and participated in the garden. Over 100,000 people have visited and looked around in the garden during that time. One of the things that surprised us was how connected or attached the people became to the seeds they planted in the garden. People were sending emails to their friends in the garden saying, `Hey, can you water my plant while I am gone?', or people were becoming very protective of their plants and getting irritated when someone else planted a seed nearby them. We also see interactions between people, talking back and forth. Generally, they are talking about things like the weather, their kids, their own gardens…* "

The Telegarden installation has therefore not only pushed the boundary of technology – the functional basis for providing automated horticulture, but also challenges our understanding of what sensorial perceptions gardening entails. Analogous to the rhetorical question "Is a Tomato natural" posed by Ann Vileisis[xxxvi], the Telegarden challenges what the key ingredients to create a pastoral experience are, and in particular which role technology can play to create an experience that is "natural".

The Distributed Robotic Garden (MIT)

Unlike the Telegarden, which focuses on actuation and relies on humans for perception and decision-making, the distributed robotic garden at MIT is a first attempt of an integral system that combines sensing, computation and actuation[xxxv]. The system consists of two mobile robots that are equipped with mobile manipulators, water bottles and cameras and tend to four potted cherry tomato plants, see also Figure 2. Each plant in the system is augmented with the capability to keep track of its own state, measure the humidity of its soil, and call a robot for help if needed by integrating a small computer, humidity sensor, and communication device into its pot. Distributing sensing and computation in this way has the following advantages: instead of maintaining a central database with potentially hundred thousands of plants, and keeping track of how tall each one is and how many fruits it bears, this information could be stored right at the plant. The plant could also use this data to perform simple reasoning such as "my fruit were green two weeks ago so they must be red now" and have it updated by robots that can be recruited to this purpose. For example a plant can ask a robot to perform an inventory of its fruit status. In order to do this, the robot positions itself in front of the plant, sweeps its camera in front of it and counts the number of red and green tomatoes it sees as well as their position. This distributed architecture is important as it allows the system to be *scalable*, that is remain operational even if the number of plants are infinite, as there is no central "bottleneck" through which information need to pass and where decisions need to be made.

Figure 3: Gardening Robot tending to tomato plants at MIT's Computer Science and Artificial Intelligence Laboratory.

The system was able to perform the following functions: respond to a watering request by a pot by autonomously docking at the plant and delivering a fixed amount of water, performing an inventory consisting of location and color of fruits on a plant and storing this information on the plants database by wirelessly communicating with its pot, and harvesting a specific fruit by obtaining its approximate location from the pot's database, visually servoing to it and grasping it with its claw.

Although the system integrates sensing, computation and actuation to perform autonomous decisions, it is far from a productive system. Key challenges of the system are: limited work-space, accuracy and dexterity of the robotic arm, limited depth-perception required for manipulation, and the robustness and accuracy of the vision system used to detect fruits. Even if these problems will be solved – which is feasible albeit not necessarily leading to an economically viable solution – operation of such a system is still heavily constrained to operation on potted plants and infrastructure that allows robots to navigate between pots. Finding the right trade-off between environmental modification, needs of producers, and fraction of tasks that need to be performed by human laborers is thus another hard problem that needs to be addressed before an autonomous gardening system will find broad acceptance.

Discussion

It is worth noting that neither agriculture nor horticulture have ever been the main drivers for technological innovation. Rather, industrial methods and techniques were adopted to agriculture long after they have proven successful in other trades[iii]. This trend seems to be persistent till today. For example, automation in agriculture is enabled exclusively by the availability of small-scale and cheap computation, which in turn has been driven by the electronic spread-sheet. Similarly, disruptive actuation and sensing technologies that have the potential to bring agricultural automation to a new level are being developed in orthogonal markets, such as the gaming industry, manufacturing, or construction. Indeed, major industrial players in precision agriculture support their technical innovations by developing for the construction industry, e.g., large earth moving equipment, which provides higher margins than agriculture or horticulture. As a result of the impact of the post-industrial agricultural system on the environment, economic drivers become more and more on par with ecological ones. Indeed, concerns about the environment and personal health have accompanied almost every technological development in agriculture, comprehensively described at the example of the tomato plant and the industrial processes surrounding it by Vileisis[xxxvi]. With an increasing understanding of future challenges on our agricultural system due to population growth and limited, non-renewable resources such as oil and phosphorus that will put hard bounds on the operation of our current system, however, this trend might reverse and agriculture becoming one of the main drivers for technological development and robotic technology in particular.

Despite providing some of the required functionality under laboratory conditions, robots are still far away from performing autonomous gardening. The key challenges are: first, the number of tasks that are involved in tending to a tomato plant is much larger than the systems described here suggest. Plants require not only watering and harvesting, but also about trimming, weeding, pollinating, cleaning, and debugging. Consequently, perception problems are not limited to finding out where the ripe fruits are, but to see whether the plant is pest infested, if the ratio between stems and sprouts is too large and the plant needs more light, or if the plant lacks nutrients and has all crumply and yellow leaves, among many others. Similarly, manipulation challenges are not limited to fruit-picking or trimming, but comprise mundane tasks such as removing an entire plant down to the root or removing larvae from under a leaf. Second, while monitoring soil humidity content is important, plant needs are very subtle and potential for wrongdoing is large. Although mankind has produced a lot of procedural and anecdotal knowledge to successfully raise crops and horticultural plants, it is very hard to distill it into clear assignments (algorithms) for a robot. A good analogy to understand this challenge might be to program a computer to play chess. Given enough computing power, one can predict all possible outcomes of different moves and rank order them. Doing this only based on the rules of the game, that is avoiding moves that let you loose a figure, performs very poorly against an human adversary, however, and only adding large amounts of anecdotal knowledge – in the form of what moves Master chess player's have decided on in a particular situation – and heuristics makes for a competitive player. Gardening is similar in that just adding water on a regular basis wouldn't get a robot very far; instead the robot needs to understand the subtle clues recorded over the lifetime of the plant and how to react to them in an optimal way.

In addition to challenges that stand in the way to have robots performing gardening tasks as we know it to raise crops and artisanal plants in our offices, kitchens and living rooms, the true challenges are to do this *better* than its done today. That is we want to improve yield while saving water, nutrients and overall energy consumption. Implementing these benefits via autonomous robots is complementary to improvements developed by the agricultural industry, which include improved fertilizers, novel plant breeds that are more productive and robust, sensor-based decision making systems used in precision agriculture, and increased automation of existing farm equipment. Although these improvements have led to massive productivity increases in the last 100 years, current trends in automation, i.e., farming equipment of increasing size, continues to drive the trend to mono-cultures, which is accompanied by the environmental and societal problems stated above. Replacing large-scale farming equipment with small-scale platforms that can take care of individual plants can therefore lead to a paradigm change: by tending to individual plants, plants can be grown in poly-cultures as well as on small patches of land that are close to the consumer. At the same time, the possible saving of resources due to individual care is the upper limit that can be achieved with any precision agriculture system.

A vast amount of research is needed to implement this vision. We need to learn how to make robots understand what they are seeing; for example, what is a plant and what is not, what is a leave and what is a fruit; we need to learn how to make robots more gentle, for example, when does the robot touch a plant and how much force it is using to push it out of the way to reach a fruit; we need to learn how robots can manipulate objects that are not rigid but flexible, for example, where does the robot has to hold the plant with its one hand so that it can pull of a tomato with its other; finally, we need to better understand what a plant needs to grow optimally, but without wasting any water or nutrients. Interestingly, robot technology can help us to achieve just that: when the robot is involved in growing the plant, it can collect data such as the plants height, its color, and how much water and nutrients it actually has deposited. By this, robots could be as revolutionary in gardening and agricultural science as they have been when sequencing genes.

A major challenge that goes along with further developing autonomous robotic technologies is the loss of conventional jobs. This is indeed the case, as more capable robots might replace seasonal workers if they can be sold cheap enough. Replacing manual labor by automated solution is an ongoing process, however, which had its most devastating impact during industrialization, for example after the advent of mechanical weaving machines that put out of work tens of thousands of people alone in Massachusetts as well as with the advent of mechanical farm equipment in the 1930ies. These technological advances were correlated with an overall productivity increase[xxxvii], which is – similar to the advent of agriculture 10.500 years ago – a major driving force for cultural development, including the amelioration of education and health care.

One might argue that robot gardeners will estrange us even more from gardening and agriculture, and thus contribute to the "disconnect" of the population with food production – the tension between the machine and the garden in the sense of Leo Marx[ii]. This argument is alleviated by two observations. First, an already very high level of automation, centralized production, and processing of food, has not prevented people from raising their own crops. Second, systems such as the Telegarden have challenged what the essence of a pastoral experience is and suggest that continuous technological evolution – from centralized monocultures to environmentally sustainable decentralized poly-cultures – might allow for increasing the involvement of people in agricultural and horticultural processes. Moving food production away from rural areas back into cities – enabled by small-scale robotic systems – in fact has the potential to involve citizens on a much deeper level in terms of both proximity and transparency. What kind of gardening activities, including passive ones such as observing the growing process, is of great interest in a space exploration context where gardening, horticulture and consumption of selfgrown vegetables might contribute to an astronaut's well-being on long duration missions[xxxviii].

Domain experts are often concerned about resource trade-off problems in fully automated systems. For example, irrigation, including drip irrigation, is considered a solved problem. Why would a robot drive around with a water bottle, if it is easy enough to bury a hose in the ground? While this is a valid concern, it should not be confused with the environmental and societal benefits that arise from plant-centric care. How this care is provided is an economical trade that needs to be addressed from a system perspective. A plant should receive the actuation that it needs, based on its status that can be inferred by sensing. *How* these sensors and actuators are implemented is less important. As soon as *some* of the sensors are expensive enough that they warrant a mobile platform, this platform can also carry out other tasks that have been carried out by static infrastructure. The problem where and how sensing, computation, and actuation are performed in the system is therefore independent of the challenge to provide individual care to plants.

Conclusion

The ability to tend to individual plants in a resource-optimal, autonomous way might lead to a paradigm change in agriculture as it will allow moving from the currently dominant mono-cultures and centralized production to decentralized poly-cultures. Poly-cultures are beneficial as they are less prone to pest infections and enable sustainable practices by exploiting companion plants, while decentralized production lowers transportation cost. Although these advantages are well known, implementing decentralized poly-cultures with current technology is too labor intensive and therefore economically not viable. Providing this labor with autonomous robots requires fundamental research in perception, that is the ability for a robot to understand what it senses, decision making, that is an understanding of how a plant reacts to certain inputs as a function of its developmental stage and environment, and actuation, that is the ability to perform all the manipulations that the plant needs, and includes weeding, harvesting and trimming among others. For the first time, agriculture itself might become the technological driver of this development as post-industrialization environmental concerns might supersede economical drivers that currently dominate agricultural innovation. Although there is no evidence yet that technology will ever be able to reach a level of maturity to implement the vision described here, the systems presented in this paper demonstrate sub-systems that implement parts of the required tasks. At the same time the value proposition of a decentralized poly-culture-based agriculture system is continuously increasing due to the decline of non-renewable natural resources that are critical for the operation of the current agricultural system, namely oil and phosphorus. We conjecture that this value proposition, which is already enormous[i], together with advances of robotic systems driven by other application domains such as autonomous cars, warehousing systems, and companion robots, will eventually replace current agricultural practices by decentralized, autonomous, plant-centric agriculture and enable an utopia such as described by Gillette, Macnie or Thomas that will not only "resolve the tension that Leo Marx, among others, has deemed irresolvable: the tension between the industrial and the agrarian orders, between the machine and the garden"[xxxix], but also make as efficient use of natural resources as possible.

---

[i] J. Klein. "The full Obama interview", October, 2008. Available online at http://swampland.blogs.time.com/2008/10/23/the_full_obama_interview, last retrieved July 7, 2011.

[ii] Leo Marx

[iii] D. Fitzgerald. "Every Farm a Factory". Yale Agrarian Studies, 2003.

[iv] P. Matson, W. Parton, A. Power, and M. Swift. "Agricultural intensification and ecosystem properties". *Science* 277 (1997):504-509.

[v] C. Dimitri, A. Effland, and N. Conklin. "The 20th century transformation of U.S. agriculture and farm policy". *Electronic Information Bulletin, USDA Economic Research Service* 3 (June 2005). Available online at http://www.ers.usda.gov/publications/eib3/eib3.htm, last retrieved March 17, 2011.

[vi] C. Driscoll, D. Whitall, J. Aber, E. Boyer, M. Castro, C. Cronan, C. Goodale, P. Groffman, C. Hopkinson, K. Lambert, G. Lawrence, and S. Ollinger. "Nitrogen pollution in the northeastern United States: Sources, effects and management options". *BioScience* 53 (2003):357-374. R. Howarth, D. Anderson, J. Cloern, C. Elfring, C. Hopkinson, B. Lapointe, T. Malone, N. Marcus, K. McGlathery, A. Sharpley,

and D. Walker. "Nutrient pollution of coastal rivers, bays and seas". *Issues in Ecology* 7 (2000):1-15. R. Socolow. "Nitrogen management and the future of food: Lessons from the management of energy and carbon". *Proc. Natl. Acad. Sci. USA* 96 (1999):6001-6008.

[vii] D. Cordell, J.-O. Drangert, and S. White. "The story of phosphorus: Global food security and food for thought". *Global Environmental Change*, 19 (2009):292-305.

[viii]Environmental Protection Agency. "Major Crops Grown in the United States" (2000), available online at http://www:epa:gov/oecaagct/ag101/cropmajor.html, last retrieved July 7, 2011.

[ix] A. Srinivasan, ed. *Handbook of Precision Agriculture: Principles And Applications*. Boca Raton: CRC Press, September 2006.

[x] Examples of companion plants are crow garlic, which masks scents of the companion plant and thus provides protection of insects, Queen Anne Lace that attracts predatory insects, or Clover that fixates nitrogen. See also Trenbath, B.R. "Plant interactions in mixed cropping communities." In *Multiple Cropping*, edited by R.I. Papendick, A. Sanchez, G.B. Triplett, pages 129–169. ASA Special Publication 27 (1976). Madison: American Society of Agronomy.

[xi] K. Iagnemma and M. Buehler eds. *Special Issue on the DARPA Grand Challenge, Part 1*, Journal of Field Robotics 23(8):461-652, 2006.

[xii] C. Borst, T. Wimboeck, F. Schmidt, M. Fuchs, B. Brunner, F. Zacharias, P. Giordano, R. Konietschke, W. Sepp, S. Fuchs, C. Rink, A. Albu-Schaeffer, and G. Hirzinger. "Rollin' Justin: mobile platform with variable base", pages 2572-2573. In Proceedings of the 2009 IEEE International Conference on Robotics and Automation. Piscataway: IEEE Press.

[xiii] S. Al-Kufaishi, B. Blackmore, and H. Sourell. "The potential contribution of precision irrigation to water conservation". In Proceedings of the 5th European conference on Precision Agriculture, pages 943-950, 2005. See also Y. Kim, R. Evans, W. Iversen, and F. Pierce. "Instrumentation and control for wireless sensor network for automated irrigation". ASAE Paper No. 061105. St. Joseph, Michigan, 2006, and W. Zhang, G. Kantor, and S. Singh. "Integrated wireless sensor/actuator networks in agricultural applications". In Proc. of ACM SenSys, page 317, November 2004.

[xiv] G. Carter. "Responses of leaf spectral reflectance to plant stress". *Am. J. Bot.*, 80 (1993):239-243.

[xv] The "EasyBloom" kit is available from http://www.easybloom.com, last retrieved June 23, 2011.

[xvi] A. Jimenez, R. Ceres, and J. Pons. "A survey of computer vision methods for locating fruit on trees". *Transactions of the ASAE*, 43(6):1911-1920, 2000. C. McCarthy, N. Hancock, and S. Raine. "Applied machine vision of plants: a review with implications for field deployment in automated farming operations". *Intelligent Service Robotics*, 3(4):209-217, 2010.

[xvii] C. McCarthy, N. Hancock, S. Raine. "Automated internode length measurement of cotton plants under field conditions". *Transactions of the ASABE* 52(6):2093–2103, 2009.

[xviii] L. Williams, J. Ayars. “Grapevine water use and the crop coefficient are linear functions of the shaded area measured beneath the canopy”. *Agric For Meteorol* 132 (2008):201–211.
[xix] S. Singh, M. Bergerman, J. Cannons, B. Grocholsky, B. Hamner, G. Holguin, L. Hull, V. Jones, G. Kantor, H. Koselka, G. Li, J. Owen, J. Park, W. Shi, and J. Teza. “Comprehensive automation for specialty crops: Year 1 results and lessons learned.” *Intelligent Service Robotics*, 3(4):245-262, 2010.
[xx] K. Tanigaki, T. Fujiura, A. Akase, and J. Imagawa. “Cherry-harvesting robot.” *Computers and Electronics in Agriculture* 63(1):65-72, 2008. Special issue on bio-robotics.
[xxi] E. van Henten, J. Hemming, B. van Tuijl, J. Kornet, J. Meuleman, J. Bontsema, and E. van Os. “An autonomous robot for harvesting cucumbers in greenhouses.” *Autonomous Robots* 13(3):241-258, November 2002.
[xxii] N. Kondo, Y. Nishitsuji, P. Ling, and K. Ting. “Visual feedback guided robotic cherry tomato harvesting.” *Transactions of the American Society of Agricultural Engineers (ASAE)*, 39(6):2331-2338, 1996.
[xxiii] Y. Edan. “Design of an autonomous agricultural robot”. *Applied Intelligence* 5(1):41-50, January 1995.
[xxiv] J. Reed, S. Miles, J. Butler, M. Baldwin, and R. Noble. “Automatic mushroom harvester development”. *Journal of Agricultural Engineering Research* 78(1):15-23, 2001.
[xxv] N. Kondo, K. Ninomiya, and S. Hayashi. “A new challenge of robot for harvesting strawberry grown on table top culture”. In *Proceedings of ASAE Annual International Meeting*, paper number 053138, Tampa, Florida, July 2005.
[xxvi] J. Noordam, J. Hemming, C. van Heerde, F. Golbach, R. van Soest, E. Wekking. “Automated rose cutting in greenhouses with 3D vision and robotics: analysis of 3D vision techniques for stem detection”. *Acta Horticulturae* 691 (2005):885–892.
[xxvii] Vision Robotics Inc. Available online at http://visionrobotics.com/vrc/index.php?option=com_zoom&Itemid=26&catid=6, last retrieved July 8, 2011.
[xxviii] D. Slaughter, D. Giles, D. Downey. “Autonomous robotic weed control systems: a review. *Comput Electron Agric* 61 (2008):63–78.
[xxix] W. Bargar. “Robots in Orthopaedic Surgery: Past, Present, and Future”. *Current Orthopaedic Practice* 463 (2007):31-36.
[xxx] W. Hopkins and P. Huner. *Introduction to Plant Physiology*, 4th edition. John Wiley & Sons Inc, 2004.
[xxxi] E. Heady and J. Pesek. “A fertilizer production surface with specification of economic optima for corn grown on calcareous ida silt loam.” *Journal of Farm Economics* 36(3):466-482, 1954. J. Doll. “A comparison of annual versus average optima for fertilizer experiments”. *American Journal of Agricultural Economics* 54(2):226-233, 1972. J. Knetsch. “Moisture uncertainties and fertility response studies”. *Journal of Farm Economics* 41(1):70-76, 1959. E. Dayan, H. van Keulen, J. Jones, I. Zipori, D. Shmuel, and H. Challa. “Development, calibration and validation of greenhouse tomato growth model: I. description of the model.” Agricultural Systems 43 (1993):143-163.

[xxxii] Q. Paris. "The von Liebig hypothesis". *American Journal of Agricultural Economics* 74(4):1019-1028, 1992.
[xxxiii] J. Williams, C. Jones, J. Kiniry, and D. Spanel. "The EPIC crop growth model." *Transactions of the ASAE*, 32(2):497-511, 1989.
[xxxiv] The Telegarden website http://goldberg.berkeley.edu/garden/Ars/, last retrieved July 6, 2011.
[xxxv] N. Correll, N. Arechiga, A. Bolger, M. Bollini, B. Charrow, A. Clayton, F. Dominguez, K. Donahue, S. Dyar, L. Johnson, H. Liu, A. Patrikalakis, T. Robertson, J. Smith, D. Soltero, M. Tanner, L. White, D. Rus. "Indoor Robot Gardening: Design and Implementation". *Intelligent Service Robotics,* Special Issue on Robotics in Agriculture 3(4):219-232, 2010.
[xxxvi] A. Vileisis, "Are Tomatoes Natural?". In "The Illusory Boundary: Environment and Technology in History", Martin Reuss, Stephen H. Cutcliffe, editors, University of Virgina Press, 2010.
[xxxvii] G. Timmins, *The last shift: the decline of handloom weaving in nineteenth-century Lancashire*, p. 19, Manchester University Press, 1993. G. Mohanty, *Labor and laborers of the loom: mechanization and handloom weavers, 1780-1840*, p. 114, 2006. Boca Raton: CRC press.
[xxxviii] S. Bishop. Here to Stay: Designing for Psychological Well-Being for Long Duration Stays on Moon and Mars, in: Lunar Settlements, Haym Benaroya, editor, CRC Press, 2009.
[xxxix] H. Segal. "Technological Utopianism in American Culture", Syracuse University Press, 1985.