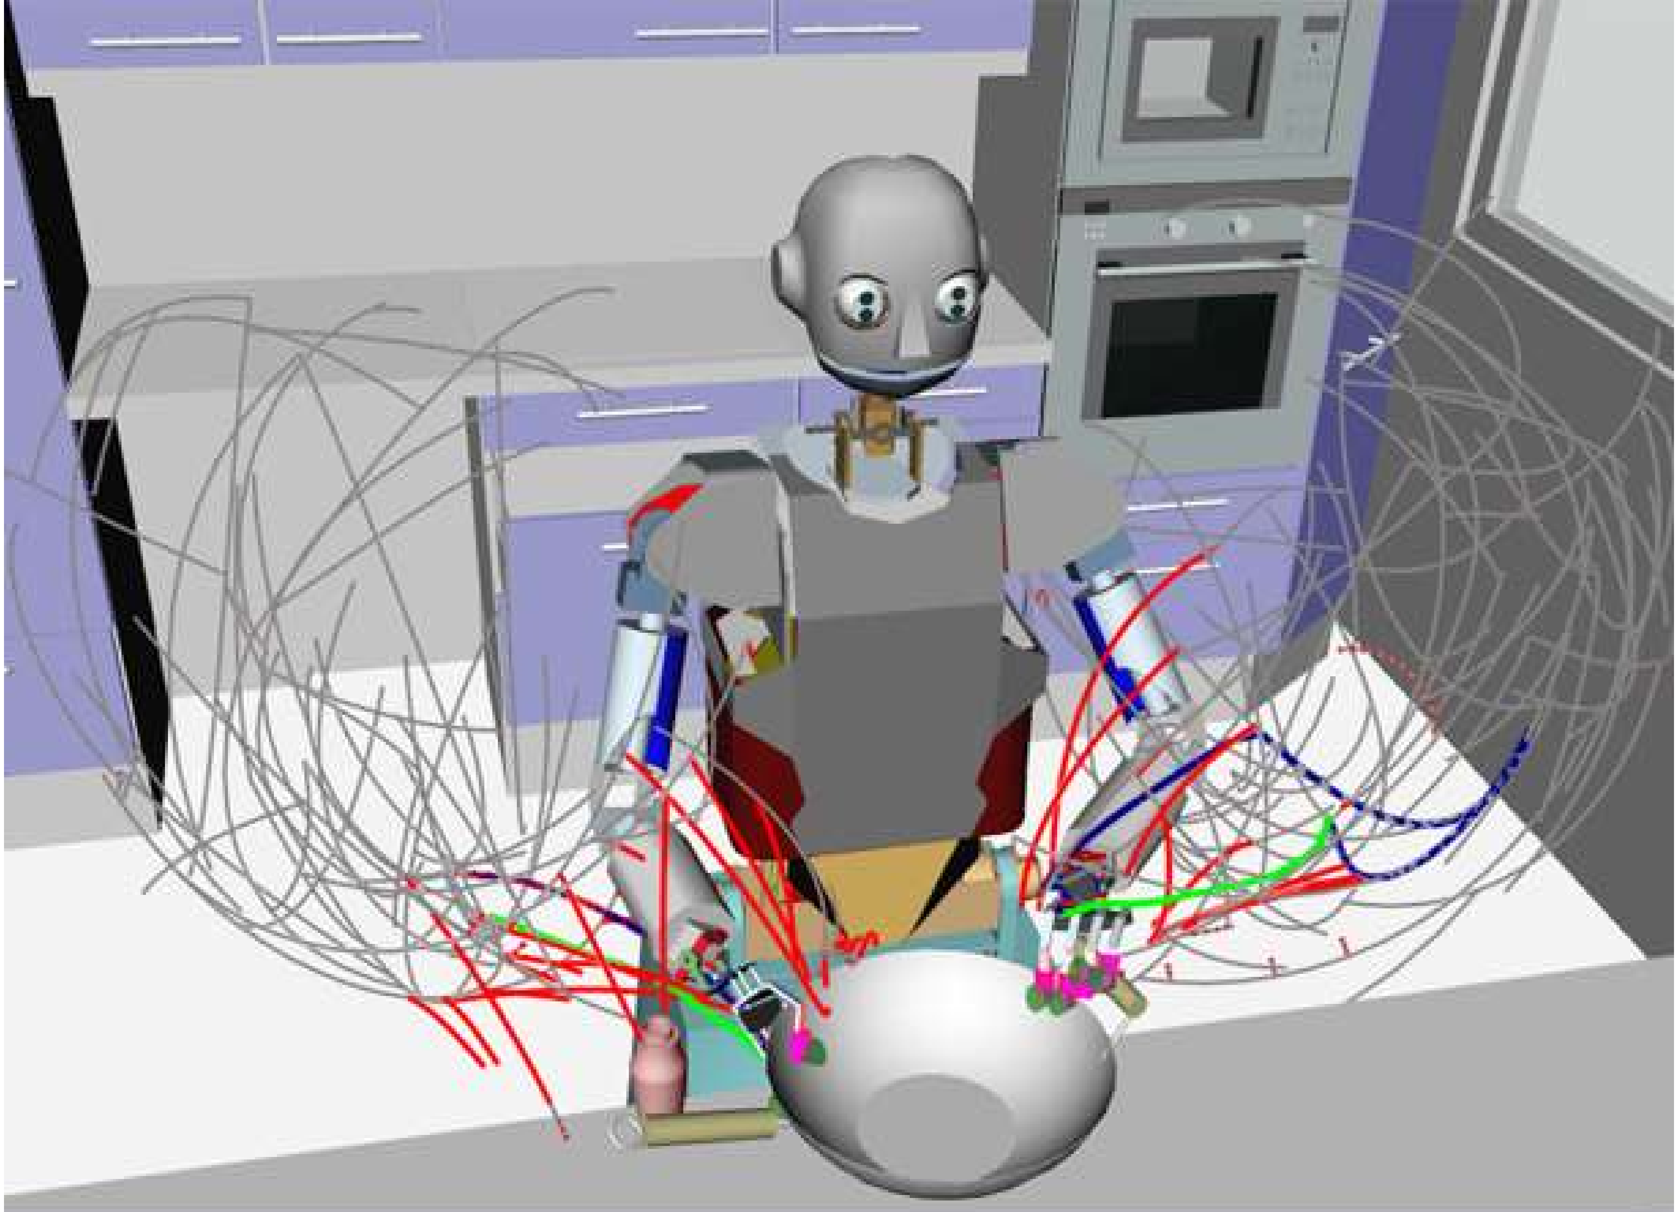

Motion Planning

(Concepts and Applications)

Leeanne Borges

First Edition, 2012

ISBN 978-81-323-3630-3

Published by:
University Publications
4735/22 Prakashdeep Bldg,
Ansari Road, Darya Ganj,
Delhi - 110002
Email: info@wtbooks.com

Table of Contents

Chapter- 1

Introduction to Motion Planning

Motion planning (a.k.a., the "navigation problem", the "piano mover's problem") is a term used in robotics for the process of detailing a task into discrete motions.

For example, consider navigating a mobile robot inside a building to a distant waypoint. It should execute this task while avoiding walls and not falling down stairs. A motion planning algorithm would take a description of these tasks as input, and produce the speed and turning commands sent to the robot's wheels. Motion planning algorithms might address robots with a larger number of joints (e.g., industrial manipulators), more complex tasks (e.g. manipulation of objects), different constraints (e.g., a car that can only drive forward), and uncertainty (e.g. imperfect models of the environment or robot).

Motion planning has several robotics applications, such as autonomy, automation, and robot design in CAD software, as well as applications in other fields, such as animating digital characters, architectural design, robotic surgery, and the study of biological molecules.

Concepts

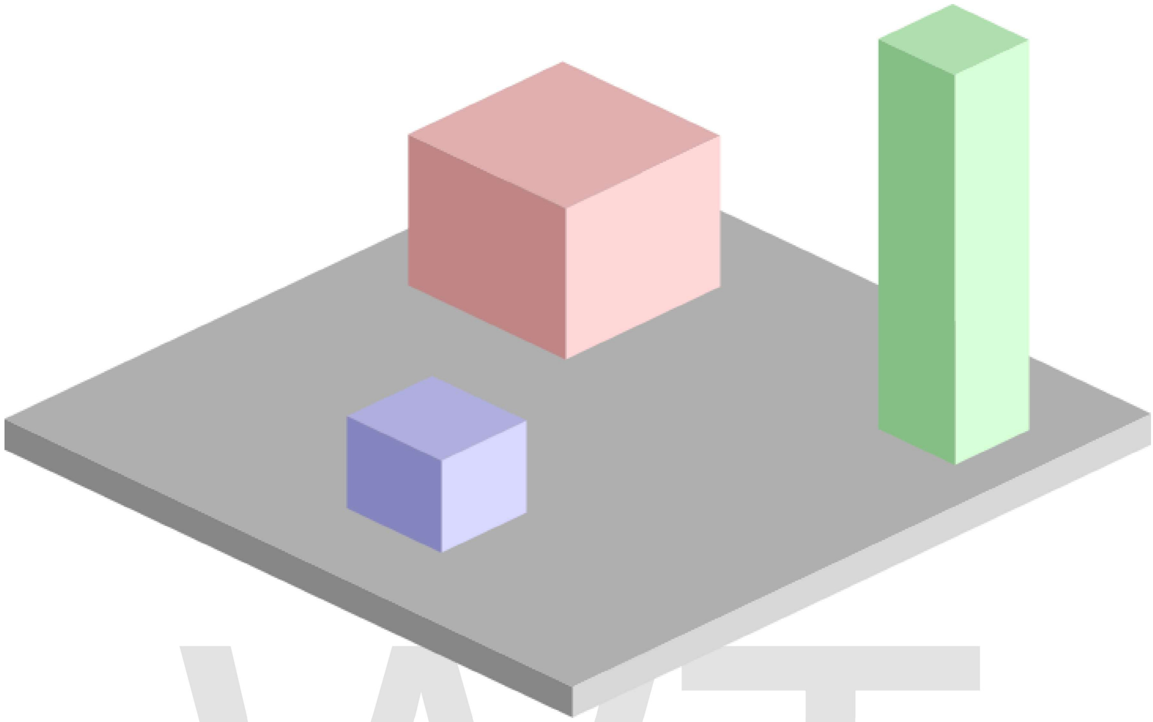

Example of a workspace.

Configuration space of a point-sized robot. White = C_{free}, gray = C_{obs}.

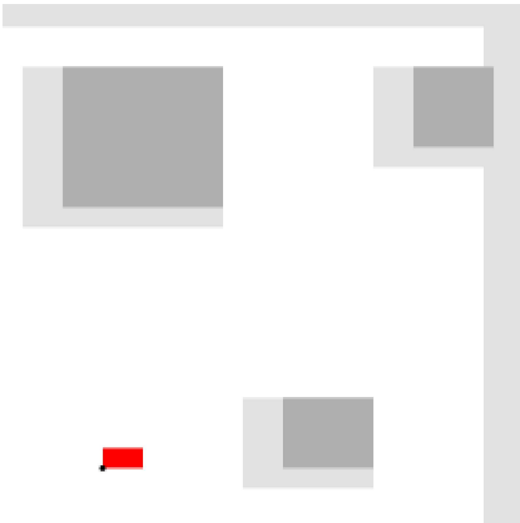

Configuration space for a rectangular translating robot (pictured red). White = C_{free}, gray = C_{obs}, where dark gray = the objects, light gray = configurations where the robot would touch an object or leave the workspace.

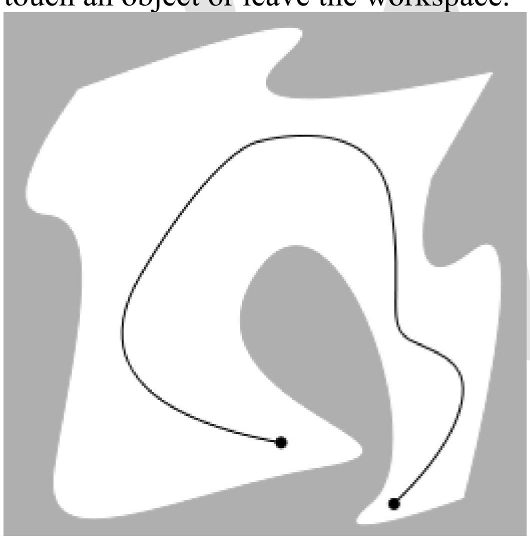

Example of a valid path.

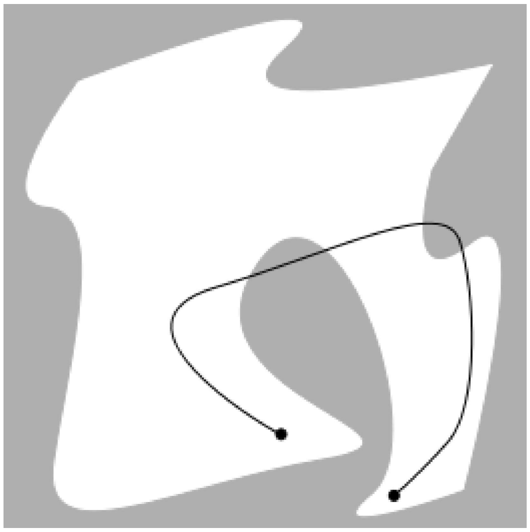

Example of an invalid path.

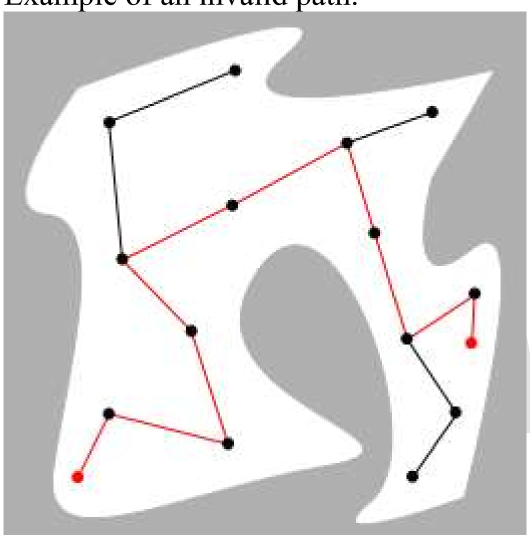

Example of a road map.

A basic motion planning problem is to produce a continuous motion that connects a start configuration S and a goal configuration G, while avoiding collision with known obstacles. The robot and obstacle geometry is described in a 2D or 3D *workspace*, while the motion is represented as a path in (possibly higher-dimensional) configuration space.

Configuration Space

A configuration describes the pose of the robot, and the configuration space C is the set of all possible configurations. For example:

- If the robot is a single point (zero-sized) translating in a 2-dimensional plane (the workspace), C is a plane, and a configuration can be represented using two parameters (x, y).
- If the robot is a 2D shape that can translate and rotate, the workspace is still 2-dimensional. However, C is the special Euclidean group **SE**(2) = $\mathbf{R}^2 \times$ **SO**(2) (where **SO**(2) is the special orthogonal group of 2D rotations), and a configuration can be represented using 3 parameters (x, y, θ).
- If the robot is solid 3D shape that can translate and rotate, the workspace is 3-dimensional, but C is the special Euclidean group **SE(3)** = $\mathbf{R}^3 \times$ **SO**(3), and a configuration requires 6 parameters: (x, y, z) for translation, and Euler angles (α, β, γ).
- If the robot is a fixed-base manipulator with N revolute joints (and no closed-loops), C is N-dimensional.

Free Space

The set of configurations that avoids collision with obstacles is called the free space C_{free}. The complement of C_{free} in C is called the obstacle or forbidden region.

Often, it is prohibitively difficult to explicitly compute the shape of C_{free}. However, testing whether a given configuration is in C_{free} is efficient. First, forward kinematics determine the position of the robot's geometry, and collision detection tests if the robot's geometry collides with the environment's geometry.

Algorithms

Low-dimensional problems can be solved with grid-based algorithms that overlay a grid on top of configuration space, or geometric algorithms that compute the shape and connectivity of C_{free}.

Exact motion planning for high-dimensional systems under complex constraints is computationally intractable. Potential-field algorithms are efficient, but fall prey to local minima (an exception is the harmonic potential fields). Sampling-based algorithms avoid the problem of local minima, and solve many problems quite quickly. They are unable to determine that no path exists, but they have a probability of failure that decreases to zero as more time is spent.

Sampling-based algorithms are currently considered state-of-the-art for motion planning in high-dimensional spaces, and have been applied to problems which have dozens or even hundreds of dimensions (robotic manipulators, biological molecules, animated digital characters, and legged robots).

Grid-Based Search

Grid-based approaches overlay a grid on configuration space, and assume each configuration is identified with a grid point. At each grid point, the robot is allowed to move to adjacent grid points as long as the line between them is completely contained within C_{free} (this is tested with collision detection). This discretizes the set of actions, and search algorithms (like A*) are used to find a path from the start to the goal.

These approaches require setting a grid resolution. Search is faster with coarser grids, but the algorithm will fail to find paths through narrow portions of C_{free}. Furthermore, the number of points on the grid grows exponentially in the configuration space dimension, which make them inappropriate for high-dimensional problems.

Traditional grid-based approaches produce paths whose heading changes are constrained to multiples of a given base angle, often resulting in suboptimal paths. Any-angle path planning approaches find shorter paths by propagating information along grid edges (to search fast) without constraining their paths to grid edges (to find short paths).

Grid-based approaches often need to search repeatedly, for example, when the knowledge of the robot about the configuration space changes or the configuration space itself changes during path following. Incremental heuristic search algorithms replan fast by using experience with the previous similar path-planning problems to speed up their search for the current one.

Geometric Algorithms

Point robots among polygonal obstacles

- Visibility graph
- Cell decomposition

Translating objects among obstacles

- Minkowski sum

Potential Fields

One approach is to treat the robot's configuration as a point in a potential field that combines attraction to the goal, and repulsion from obstacles. The resulting trajectory is output as the path. This approach has advantages in that the trajectory is produced with little computation. However, they can become trapped in local minima of the potential field, and fail to find a path.

Sampling-Based Algorithms

Sampling-based algorithms represent the configuration space with a roadmap of sampled configurations. A basic algorithm samples N configurations in C, and retains those in C_{free} to use as *milestones*. A roadmap is then constructed that connects two milestones P and Q if the line segment PQ is completely in C_{free}. Again, collision detection is used to test inclusion in C_{free}. To find a path that connects S and G, they are added to the roadmap. If a path in the roadmap links S and G , the planner succeeds, and returns that path. If not, the reason is not definitive: either there is no path in C_{free}, or the planner did not sample enough milestones.

These algorithms work well for high-dimensional configuration spaces, because unlike combinatorial algorithms, their running time is not (explicitly) exponentially dependent on the dimension of C. They are also (generally) substantially easier to implement. They are probabilistically complete, meaning the probability that they will produce a solution approaches 1 as more time is spent. However, they cannot determine if no solution exists.

Given basic *visibility* conditions on C_{free}, it has been proven that as the number of configurations N grows higher, the probability that the above algorithm finds a solution approaches 1 exponentially . Visibility is not explicitly dependent on the dimension of C; it is possible to have a high-dimensional space with "good" visibility or a low dimensional space with "poor" visibility. The experimental success of sample-based methods suggests that most commonly seen spaces have good visibility.

There are many variants of this basic scheme:

- It is typically much faster to only test segments between nearby pairs of milestones, rather than all pairs.
- Nonuniform sampling distributions attempt to place more milestones in areas that improve the connectivity of the roadmap.
- Quasirandom samples typically produce a better covering of configuration space than pseudorandom ones, though some recent work argues that the effect of the source of randomness is minimal compared to the effect of the sampling distribution.
- If only one or a few planning queries are needed, it is not always necessary to construct a roadmap of the entire space. Tree-growing variants are typically faster for this case (single-query planning). Roadmaps are still useful if many queries are to be made on the same space (multi-query planning)

Completeness and Performance

A motion planner is said to be complete if the planner always produces a feasible path, when one exists. Most complete algorithms are geometry-based. The performance of a complete planner is assessed by its computational complexity.

Resolution completeness is the property that the planner is guaranteed to find a path if the resolution of an underlying grid is fine enough. Most resolution complete planners are grid-based. The computational complexity of resolution complete planners is dependent on the number of points in the underlying grid, which is $O(1/h^d)$, where h is the resolution (the length of one side of a grid cell) and d is the configuration space dimension.

Probabilistic completeness is the property that as more "work" is performed, the probability that the planner fails to find a path, if one exists, asymptotically approaches zero. Several sample-based methods are probabilistically complete. The performance of a probabilistically complete planner is measured by the rate of convergence.

Incomplete planners do not always produce a feasible path when one exists. Sometimes incomplete planners do work well in practice.

Automation

KUKA Industrial Robots being used at a bakery for food production

Automation is the use of control systems and information technologies to reduce the need for human work in the production of goods and services. In the scope of industrialization, automation is a step beyond mechanization. Whereas mechanization provided human operators with machinery to assist them with the muscular requirements of work, automation greatly decreases the need for human sensory and mental requirements as well. Automation plays an increasingly important role in the world economy and in daily experience.

Automation has had a notable impact in a wide range of industries beyond manufacturing (where it began). Once-ubiquitous telephone operators have been replaced largely by automated telephone switchboards and answering machines. Medical processes such as primary screening in electrocardiography or radiography and laboratory analysis of human genes, sera, cells, and tissues are carried out at much greater speed and accuracy by automated systems. Automated teller machines have reduced the need for bank visits to obtain cash and carry out transactions. In general, automation has been responsible for the shift in the world economy from industrial jobs to service jobs in the 20th and 21st centuries.

Advantages and disadvantages

The main advantages of automation are:

- Replacing human operators in tasks that involve hard physical or monotonous work.
- Replacing humans in tasks done in dangerous environments (i.e. fire, space, volcanoes, nuclear facilities, underwater, etc.)
- Performing tasks that are beyond human capabilities of size, weight, speed, endurance, etc.
- Economy improvement. Automation may improve in economy of enterprises, society or most of humanity. For example, when an enterprise invests in automation, technology recovers its investment; or when a state or country increases its income due to automation like Germany or Japan in the 20th Century.

The main disadvantages of automation are:

- Technology limits. Current technology is unable to automate all the desired tasks.
- Unpredictable development costs. The research and development cost of automating a process may exceed the cost saved by the automation itself.
- High initial cost. The automation of a new product or plant requires a huge initial investment in comparison with the unit cost of the product, although the cost of automation is spread in many product batches.

Relationship to unemployment

Multivariate effect

Most people consider it common sense that automation has the potential to foster unemployment, because it obviates human work by transferring tasks to machines. However, the translation of that potential into observed effect has largely not happened in the two centuries during which it has been continually predicted. After many decades of automation development and dissemination, the net macroeconomic effect has been generally positive—automation has been part of a general trend of economic growth worldwide; standards of living have risen in many places; and automation has never yet been shown to have induced any widespread structural unemployment. The main explanation for this is that, so far, job losses in any one particular economic niche have always been more than offset by job gains in other niches. As the lowered unit cost of goods and services (which the automation made possible) gave consumers more purchasing power to devote to other goods and services, new jobs sprang up in the production of those goods and services. Thus each time that automation has freed up human resources, those resources have been redeployed by market forces (although it did not always happen without turbulence in the lives of individual workers).

One of the earliest promises of automation was to allow more free time, without any threat of income reduction. This effect has been seen in many individual facets of life (for example, the automatic washing machine has made laundry less time-consuming; engine control units have reduced the amount of automotive downtime; the automatic dishwasher has made dishwashing less time-consuming), but the net outcome of modern life in developed economies remains a state of hurry and busyness, mostly because rising living standards have brought rising expectations in direct relation. (Each time-saving improvement has made room for a new aspiration to take its place.)

Automation also does not imply unemployment when it makes possible tasks that were unimaginable without it (such as exploring Mars with the Sojourner rover). Likewise with fields where the economy is already fully adapted to an automated technology, and the jobs were lost long enough ago that the displacement was long since absorbed by the workforce (as with the continually advancing automation of the telephone switchboard, which eliminated most telephone operator jobs and kept many more from ever existing in the first place).

Today automation is quite advanced (relative to just a few lifetimes ago), and it continues to advance with an accelerating pace throughout the world. Although it has been encroaching on ever more skilled jobs, the general well-being and quality of life of most people in the world (where political factors have not muddied the picture) have improved. Clearly a multivariate effect has been at work (something much more than just the obvious idea that automation has the *potential* to cause unemployment). In fact, the idea that automation posed an *imminent* threat to employment, first articulated in 1811 by a group of textile workers known as Luddites, has proven to be so fallacious over the ensuing two centuries that economists call the imminent-threat idea the *Luddite fallacy*. Today the eternal fallaciousness of the Luddite premise is a mostly undisputed principle of economic theory, because it has proven true empirically time and again.

There is some concern today that the economy's ability to continue absorbing ever-increasing automation without experiencing significant structural unemployment may be heading toward an upper limit—that is, that we are approaching a point where the Luddite premise will no longer be entirely fallacious, because the relationship of humans to machines that made it fallacious is changing. In this view, the empirical strength of the eternal-fallaciousness idea is only a reflection of the parameter values of the environment thus far. In other words, the idea is undoubtedly an excellent explanation of the past, but whether it can accurately predict the future is an independent problem. Like an investment prospectus, proponents of this view caution that "past performance is no guarantee of future results."

Timeline of concerns about automation's relationship to unemployment

Early in the Industrial Revolution

Historical concerns about the effects of automation date back to the very beginning of the Industrial Revolution, when a social movement of English textile machine operators in the early 19th century known as the Luddites protested against Jacquard's automated weaving looms. The Luddites destroyed a number of these machines, which they felt threatened their jobs.

Later in the Industrial Revolution

The development of the American system of manufacturing disgusted many skilled machinists at a time when the very definition of being a machinist included a core element of skilled toolmaking and fitting on a craft basis. Innovations of this system included increasing reliance on jigs and gauges and on machine tools that built more of a process into the tool's movements (such as turret lathes and screw machines). These innovations continually turned skilled work into semi-skilled or unskilled, contributing to vast migrations of laborers across borders and oceans. However, despite this transformation, there were always other economic niches for skilled workers to go to, given enough searching. Recessions interfered with employment, but no foundational aspects of structural unemployment were caused by automation itself.

During the Machine Age

As in the preceding century, the period of 1880 to 1940 saw no underlying automation-induced structural lack of new economic opportunities for skilled workers to go to, given enough searching, although the Great Depression caused a tremendous disruption to employment. The foundational *potential* for full employment had not been lost, as would later be shown by the post–World War II economic expansion and other economic miracles.

During the 1950s through 1990s

The postwar development of new automation technologies using electronics, servomechanisms, and digital computers stoked a new wave of fears similar to the old Luddite ones. Among the working class and labor unions, there was stiff resistance to loss of employment through automation, including contract clauses won in hard-fought contract negotiations that mandated alternate employment for any workers whose positions were eliminated by automation. These clauses seemed a great victory for union workers at large corporations in developed nations, but because they had no effect at smaller, nonunionized companies or in developing nations, those corporations faced withering competition that shrank their market shares until their workers' gains eventually undermined their own success. However, the salvation for employment rates damaged in the industrial sector (secondary sector of the economy) came from the service sector (tertiary sector), which absorbed all of the workers that automation displaced elsewhere. For example, many manufacturing jobs left the United States during the 1990s but were offset by a one-time massive increase in IT jobs at the same time. And in some cases the freeing up of the labor force allowed more people to enter higher skilled managerial jobs and technically specialized jobs, which are typically higher paying. Therefore, fears of unemployment due to automation were generally dismissed as just another instance of the Luddite premise, which had proven fallacious time and again over many decades. Given this obvious empirical contradiction of the premise, people who nevertheless returned to it were usually viewed by the mainstream as cranks misled by quixotic leftist political bias. For example, works by scholars including David F. Noble and Jeremy Rifkin were often respected but discounted. At worst, they were mocked with the disparaging label "neo-Luddite". Noble even wrote a later book titled *Progress Without People: In Defence of Luddism* to try to further explain why the Luddite premise should not be laughed out of academia.

Post-market musings

Rifkin's *End of Work*, published in 1995 and written by a non-engineer, predicted automation-induced unemployment despite having a rather hazy idea of how IT would evolve over the next decade. (The book mentioned the Internet once in passing and the World Wide Web not at all. Its IT focus was mostly on robotics.) Also hazy was Rifkin's explanation of any solution to the problem. The book's subtitle called the solution a "post-market economy", but its concluding chapters did not clearly lay out how such an economy could be engineered, leaving readers to conclude that a non-market solution involving a planned economy was implied between the lines.

In terms of political economy implications, there was no clear differentiation at the time between the ideas of authors like Noble or Rifkin (on the one hand) and traditional leftist agitation (on the other hand). To the extent that readers could ask "What point is this guy getting to?" and answer the question with "socialism" or "a welfare state", they dismissed these authors.

During the 2000s and 2010s

Since the 1990s, the possibility has been raised again in even an apolitical, technocratic way that the Luddite premise (that automation creates unemployment) was only fallacious in the absence of highly advanced and ubiquitous automation, which until recently was mostly out of reach technologically. This would explain why it has always been fallacious until now, but also why it might not always remain so. For example, Marshall Brain, Martin Ford, and others have suggested that exponentially accelerating information technology (IT) may ultimately result in widespread structural unemployment, because an implicit assumption underlying the "eternally fallacious" idea (that lots of regular humans will always find ways to do service work that machines can't do) will itself be fallacious as IT advances. They suggest that, unlike in the 20th century, when the tertiary sector absorbed all of the workers that the automation of the secondary sector expelled, the tertiary sector now also faces depopulation via automation; its employment will shrink, not grow, and this time there is no other sector to backstop the process by absorbing the displaced workers. The high unemployment rates of the late-2000s recession have brought the idea of structural unemployment back into mainstream attention, as observations are made about positions that require extensive specialized skill and experience standing long vacant even while general unemployment rates above 9% (and horror stories of fruitless job searches) would seem to suggest that such vacancies ought to be scarcer. The idea that automation has finally advanced to the point that the Luddite premise is no longer entirely fallacious is one of the components of some theoretical explanations for the string of jobless recoveries in developed economies in recent decades. Expectations that the (already eroding) fallaciousness will fall off sharply in coming decades underlie the fear of structural shift.

Writers such as Rifkin, Brain, and Ford often suggest that the structure of the economy will have to shift to a basic income because its present structural foundation (trading labor for income) will no longer be an available option. They often include an element of civic obligation, such that able people must somehow contribute civically in order to receive the basic income. The labor-market economy (trading labor for income) already achieves that outcome today (because working for income generally produces civic value in various ways, directly and indirectly), but the argument is that advanced automation will decouple the linkage that makes that possible. Thus the same result (trading civic value for income) would have to be driven by different forces—either non-market ones, or via a new kind of market. The non-market idea seems infeasible given the generally abysmal performance record of planned economies. But the idea of engineered new markets leaves room for the disciplining and motivating powers that make capitalist markets capable of positively shaping human behavior where government alone is usually unable.

New-market engineering

Brain and Ford's books, in stark contrast to Rifkin's, came later and were written by engineers with extensive under-the-hood knowledge of modern production methods, computer hardware and software, and the Internet's underpinnings. They explicitly reject non-market solutions as unworkable and instead suggest new kinds of markets. Rather than being "post-market" proponents, such authors could be called "new-market" proponents. They vigorously distance themselves from socialism or welfare states—generally seeking to keep a market economy with private enterprise, which they believe cannot be preserved *unless* its foundation is modified from its current structure. Thus, quite contrary to being anti-market agents (as critics might suppose them to be), they believe themselves to be *salvaging* markets from destruction. They envision creating consumer purchasing power by some other mechanism than the traditional labor market as we have known it so far, in order that free markets may continue to provide the invisible hand component of production-possibilities decisions. In other words, they believe that market forces are necessary to generate allocative efficiency, and they believe that without a structural modification that (at least partially) decouples purchasing power (and consumer confidence) from employment determined by the traditional labor market, there will be a systemic market failure, which they seek to avoid.

Just as new-market advocates are pro-market and pro-private-property, they are also very much non-Luddite (in fact, exactly opposite of Luddite) in the respect that they *like* technology—they don't *hate* it. They want it to continue advancing as robustly as ever. They simply feel that income and purchasing power must be decoupled from human participation in production. (The decoupling does not have to happen all at once; it could start small and gradually increase.) If that happens, then they essentially do not have any problem with technology or automation, per se. In contrast to old-style welfare, they do not feel that income should be unconditional, or equal, or "free" (given out "for nothing"). They believe that people should have to work for it (in a new sense of the word "work"), in the respect that they are given incentives to do positive things, like take classes, read books, conserve environmental resources, and so on. People would be paid to do civically valuable things, and if they chose not to do those things, they would not be paid. In this way, new-market advocates align themselves with human nature, which generally requires selfish motivations and incentives to shape behavior, and with the market's invisible hand, which is needed to make the right production-possibilities decisions (because the idea that individual human managers, or groups of them, are capable of making those decisions correctly with zero invisible-hand assistance has been empirically discredited).

Wage-recapture market variant

Ford's main new-market mechanism would be to create a tax that recaptures most (not all) of the value that firms and their customers gain from eliminating wages, then use the tax revenue to pay people for doing civically valuable actions—that is, pursuing activities, such as higher education or environmental preservation, that have positive externalities. The main reason for paying these "wages" need not be their altruistic or environmentalist components; the main reason is simply to prevent the market economy from collapsing due to noncirculation of value (that is, the lack of adequate trade which would occur if lack of consumer purchasing power and confidence left no way for an adequate mass market to exist). Ford points out that the tax could not take *all* of the gains away from the corporations and their customers, because this would destroy the natural incentive to innovate that a market economy needs to be sustainable. The value would be split between the innovators, their customers, and the rest of the population, because leaving out any of that trio would wreck the sustainability of the model. (The leaving out of the third leg is what is causing today's economic pathologies and promising tomorrow's, in the view of new-market engineers.) Ford's idea is an earnest market effort because it preserves the invisible hand as the maker of production-possibilities decisions for goods and services. However, it does rely on human planning (via a technocratic government agency in each country) to make the production-possibilities decisions for civic actions. The latter is viewed as unfortunate but necessary due to the lack of an alternative.

Mirror-image market variant

Another idea for a new-market mainspring which solves the aforementioned "lack of an alternative" problem is a "mirror image" idea, which has an even more private-sector approach in which the invisible hand helps make even the civic-actions production-possibilities decisions. In this model, the government does not collect a wage-recapture tax at all. Instead of enforcing tax payment, it only enforces payment of a new-style "wage" directly from corporations to consumers that looks to us today like something we might label "mandatory philanthropy", but which would actually be a true market wage of a new type. In today's old market, money flows from consumers, through (partially automated) companies, past the eyes of the government enforcement sentry (but not through its hands) as wages, into the hands of workers (who are also the consumers, thus completing the cycle of value recirculation). In the new market (mirror-image variant), money would flow from consumers, through (highly automated) companies, past the eyes of the government enforcement sentry (but not through its hands) as [new-style] "wages", and into the hands of [new-style] "wage" earners, who are paid the "wage" for civically valuable actions. (They are also the consumers, thus completing the cycle of value recirculation).

In this model, the decisions about what the civic actions are can be made by the invisible hand, because each "mandated philanthropist" gets a large degree of authority in what actions their "philanthropy" (which is actually [functionally] a new-style "payroll") will or won't pay for. Many such paymasters functioning simultaneously could constitute the "buyers" in a market for civically valuable actions (with mass-market "workers" as the "sellers"). There would still be *some* regulation involved, because, for example, it would be illegal to base the "payroll" decisions on race, color, religion, creed, gender, sexual orientation, ethnicity, disability, marital or veteran status, and so forth. To decide which "workers" were on a given "payroll", there might be a clearinghouse to randomly match the two, rotating assignments every several years. Or perhaps the businesses that run the "payroll" could even "hire" the "workers" themselves, in which case workers would compete for "jobs" by showing off how "productive" they could be in doing the civic actions (another level of invisible hand yet again). The "mirror image" name comes from the idea that this variant of new market is a very free market where the invisible hand remains just as powerful as it was in the 1945-2008 economy, but with many mirror-image aspects (which are visualized above by the amount of quotation marks that are necessary to signify mirror-image senses for words that were always [up till now] widely known only in their non-mirror-image senses).

The axis of reflection in the mirroring seems to be, at root, a "polarity shift" from where human individuals can add value only by *doing production* (from within production systems) to where they can also add value by *avoiding hurting production systems* (from outside). The hurt-avoidance comes from such civic actions as providing goods-and-services demand via consumption (which the system requires in order to stay running) instead of failing to consume (because of lack of income); ensuring the sustainable supply of energy and environmental resources to the production systems (by avoiding *over*consuming those); and by ensuring the supply of people educated enough to provide the few humans that the production systems will need in the future, by pursuing education and cognitively enriching pastimes. The humans that the systems need will be few, but those few will need to be highly intelligent, talented, and educated, constituting a human resource that might be endangered if the general population does not act as a "farm team system" for it by valuing education and self-education as a civic action. An analogy is provided by sports' relationship to general life. Few humans are talented and practiced enough to play professional sports, but the professional teams rely on a system that filters such scarce people out of the general population via little league/pee wee programs, high school play, college play, farm-team play, etc. People in the general population are not considered inferior human beings (versus the pro players) because of their lack of pro talent. They are valued as the fans and ticket-buyers that make the pro system economically viable. And a small fraction of them grow up to become pros themselves.

In today's old market, governments enforce the payment of wages by having outlawed their nonpayment (i.e., slavery); by levying tariffs on cheap competition from countries that kept their nonpayment (slavery) (that outlawing has now been global for many decades); and by attempting to minimize their underpayment (i.e., wage slavery, a sharply cheapened value of work [with elites and their customers keeping the money]). In the new market (mirror-image variant), governments enforce the payment of "wages" by outlawing their nonpayment (i.e., evading the "payroll"); by levying tariffs on cheap competition from countries that kept their nonpayment (non-participating countries); and by attempting to minimize their underpayment (i.e., "wage" slavery, a sharply cheapened value of civic actions [with elites and their customers keeping the money]).

One of the inherent challenges of the mirror-image variant is that various forms of dressing up corporations' financial self-interest in a specious cloak of civic virtue would inevitably arise. This would be a "washing" form of marketing and operations that included greenwashing and analogous washing in other domains of life (e.g., education, infrastructure). It seems unlikely that this can be entirely negated; instead, it would have to be perennially pruned by social censure and regulatory oversight. However, no other system is without its chronic weaknesses, either. For example, the classical variants of capitalism (implemented thus far) have scored poorly on various tests, such as environmental sustainability and (potentially) the employability of the average human (as that was traditionally defined) as automation grows pervasive. Twentieth-century variants of communism fared even worse in environmental sustainability, and also failed economically in average standard of living and politically in individual freedom. The wage-recapture new-market variant, with its technocratic decisions on how to spend the revenue, holds promise to minimize the corporate "washing" problem, yet it also holds risks of failing on allocative efficiency and market-driven innovation, which the mirror-image variant mitigates. As elsewhere in reality, each choice has pros and cons, rather than any choice being perfect. The "washing" problem may be the mirror-image analog of classical capitalism's tendency to exaggerate needs (for example, a maker of antibacterial soaps encouraging the populace to fear microbes to an irrational degree). Both are forms of conflict of interest that cause "chronic irritation" to a socioeconomic system but need not be "fatal" to it if given adequate "medical management". The washing problem may be less systemically injurious than the allocative inefficiency problem, just as the "exaggerated needs" problem of classical capitalism was less systemically injurious than the allocative inefficiency problem of central economic planning.

In choosing the decider of production-possibilities decisions (whether of goods, services, or civic actions), the invisible hand is generally preferred to committees of humans because it has proven to be superior at the decision making (except for regulatory issues such as race-color-religion-etc and the "washing" discussed above). In the future it will also be necessary to ask what role artificial intelligence might possibly have in making those decisions, and whether humans would allow it. Perhaps artificial intelligence, like human intelligence, will share the role with the invisible hand but be barred from usurping the entirety of it.

Implementations

Regarding the chances of any new-market ideas being implemented, there are both significant barriers and significant drivers, with a net potential of perhaps "even chances". Ford discusses many of these barriers and drivers. The barrier side includes (a) natural cultural conservatism that powerfully resists systemic changes; (b) the powerful influence of laissez-faire ideals, which would resist any engineered systemic change to markets (especially *anything* involving a tax); (c) the fact that early implementation by individual countries faces an immediate threat from the export and offshoring competition of countries that *haven't* yet implemented; and (relatedly) (d) the all-or-nothing problem, which may occur if a new system would work well but only if the switch from old to new was an off-on switching rather than a gradual evolution. However, on the driver side there are powerful forces that may answer all of the barriers. Foremost would be a dawning realization by economic elites that they have a choice between a new market with prosperity, or the old market spiraling into near-total failure. Globalization so far has not threatened the wallets of economic elites (only those of average workers), and has in fact enriched the elites thus far; but the changing parameter values of the economic system as automation advances would alter that runtime environment and transform it into a new one, where even the elites' wealth would be threatened by a market failure that killed their businesses and reduced asset values throughout the economy. Realizing these options, elites might actually switch from opposing new markets to actively supporting their implementation (including addressing the competition between countries whose policies differed). The all-or-nothing problem does not have to occur if an implementation is engineered such that extremely profitable, extremely automated industries began piloting new markets while other industries continued with an old-market status quo for quite some time. In this model, the early adopters voluntarily become leaders, and the pilot projects would act simply as economic stimulus on the broader economy (although a type of stimulus much more effective than old-style stimulus, whose efficacy seems to be eroding because it relies on the Luddite premise being a total fallacy as opposed to shifting by degrees out of total fallaciousness). The overall transition in this model could actually be quite painless, as a generally prosperous economy changed gradually over decades from mostly-old-with-some-new to mostly-new-with-some-old. The selfish motivation of the early-adopter leaders would be the aforementioned choice faced by economic elites. They would choose to stimulate the broader economy because that result would ensure their own continuing strong sales and growth by preserving a runtime environment of general prosperity for them to operate within, without which depression or malaise would occur.

Given the aforementioned choice faced by economic elites, those in the private sector might even choose to pursue the new market without government involvement. But the private sector faces two hurdles that would make it difficult: the natural competition between firms (which is necessary and thus must be protected by competition law), and legal obligation to maximize shareholder value. The traditional definitions of shareholder value evolved in an earlier era whose commercial environment had different parameter values due to lack of advanced automation. Those traditional definitions would bar new-style payrolls. But in the face of market failure without them, perhaps a case would emerge for an updated definition. Competition is the other hurdle. Companies are barred by competition law from agreeing to limit competition, and even if they weren't, individual companies generally cannot make the first move of increasing expense without being killed by competition from rivals who don't. This is why "a level playing field" would have to be created by policy, or to use a different analogy, "a high tide that lifted all boats equally". This is directly analogous to existing minimum wage laws. Individual companies generally could not survive in the market if they volunteered to self-enforce minimum limits on wages (in the absence of any laws requiring them). There *is* breathing room for above-market wages (e.g., to attract superior talent) at some companies in some industries who enjoy a relatively high level of imperfect competition; but most companies in most industries face competition too close to pure to survive the attempt. In this sense, the mandated value recirculation (whatever anyone calls it, from "wage recapture" to "new-style wages") is as unremarkable and non-novel an idea as any legislative or regulatory mechanism in commerce. For goals that make long-term systemic balance possible but cannot be pursued by the self-interest of individual market players, these mechanisms provide a path by forcing all competitors to play the game by the same rules. Existing examples include employment standards (e.g., child labor laws, minimum wage laws), environmental protection, and financial regulation (to prevent bubbles and thus crashes). These exist in perennial tension with the forces of pure capitalism; thus the extremes perform checks and balances on each other. Businesses usually fight for inadequate regulation; government usually fights for excessive regulation; and a sustainable balance results. Over decades, systemic pathologies gradually push the balance point out of the sustainable range; periodic breakdowns then yield correction by counteractive forces (e.g., trust-busting [leftward correction], the Reagan revolution [rightward correction]).

Laissez-faire ideals reigned supreme worldwide for about three decades (roughly centered on the fall of the Soviet Bloc, which vindicated capitalism over central planning in many ways). In this environment, where the lesson commonly extrapolated was that pure capitalism will always be better than any mixed-economy alternatives, the prevailing theory has been that higher corporate taxes can only harm economic prosperity. The reasoning is partly that countries can simply compete to undercut each other's corporate tax rates (which is true), but also, more importantly, that only the invisible hand is capable of recirculating capital back toward the base of the economy in a successful manner (which is not to be dismissed lightly, and may in fact be true). The disparaging label for such ideas is "trickle-down economics", but many intelligent people have earnestly believed in these ideals; and the fact that their discounting has often been facile and done by imperfect opponents has only encouraged believers to stay faithful. Widespread fervor for trickle-down beliefs (in both the public and private sectors) poses a formidable barrier to the wage-recapture-tax new-market variant. But these conventional beliefs rely on the assumption that the Luddite premise is entirely and eternally fallacious. Unfortunately, there has already been a decade of empirical evidence that low taxes, new business investment, and economic growth no longer have a sure-fire correlation to strong, "good-jobs" employment in developed economies. If the Luddite premise has been starting to shift into partial accuracy, then no amount of continued low taxes and deregulation will ever be able to produce enough trickling down to create broad-based prosperity. In that case, mandating the payment of new-style wages could recirculate value back to the base of the mass market. The promise of the mirror-image variant would be that humans need not turn to central planning for the distribution details, because as long as the "minimum wage" (referring to the new-style wages) and other employment standards are being enforced, then government's role ends there.

Other goals of automation (beyond productivity gains and cost reduction)

In manufacturing, the purpose of automation has shifted to issues broader than productivity and costs.

Reliability and precision

The old focus on using automation simply to increase productivity and reduce costs was seen to be short-sighted, because it is also necessary to provide a skilled workforce who can make repairs and manage the machinery. Moreover, the initial costs of automation were high and often could not be recovered by the time entirely new manufacturing processes replaced the old. (Japan's "robot junkyards" were once world famous in the manufacturing industry.)

Automation is now often applied primarily to increase quality in the manufacturing process, where automation can increase quality substantially. For example, automobile and truck pistons used to be installed into engines manually. This is rapidly being transitioned to automated machine installation, because the error rate for manual installment was around 1-1.5%, but has been reduced to 0.00001% with automation.

Health and environment

The costs of automation to the environment are different depending on the technology, product or engine automated. There are automated engines that consume more energy resources from the Earth in comparison with previous engines and those that do the opposite too. Hazardous operations, such as oil refining, the manufacturing of industrial chemicals, and all forms of metal working, were always early contenders for automation.

Convertibility and turnaround time

Another major shift in automation is the increased demand for flexibility and convertibility in manufacturing processes. Manufacturers are increasingly demanding the ability to easily switch from manufacturing Product A to manufacturing Product B without having to completely rebuild the production lines. Flexibility and distributed processes have led to the introduction of Automated Guided Vehicles with Natural Features Navigation.

Digital electronics helped too. Former analogue-based instrumentation was replaced by digital equivalents which can be more accurate and flexible, and offer greater scope for more sophisticated configuration, parametrization and operation. This was accompanied by the fieldbus revolution which provided a networked (i.e. a single cable) means of communicating between control systems and field level instrumentation, eliminating hard-wiring.

Discrete manufacturing plants adopted these technologies fast. The more conservative process industries with their longer plant life cycles have been slower to adopt and analogue-based measurement and control still dominates. The growing use of Industrial Ethernet on the factory floor is pushing these trends still further, enabling manufacturing plants to be integrated more tightly within the enterprise, via the internet if necessary. Global competition has also increased demand for Reconfigurable Manufacturing Systems.

Automation tools

Engineers now can have numerical control over automated devices. The result has been a rapidly expanding range of applications and human activities. Computer-aided technologies (or CAx) now serve the basis for mathematical and organizational tools used to create complex systems. Notable examples of CAx include Computer-aided design (CAD software) and Computer-aided manufacturing (CAM software). The improved design, analysis, and manufacture of products enabled by CAx has been beneficial for industry.

Information technology, together with industrial machinery and processes, can assist in the design, implementation, and monitoring of control systems. One example of an industrial control system is a programmable logic controller (PLC). PLCs are specialized hardened computers which are frequently used to synchronize the flow of inputs from (physical) sensors and events with the flow of outputs to actuators and events.

Human-machine interfaces (HMI) or computer human interfaces (CHI), formerly known as *man-machine interfaces*, are usually employed to communicate with PLCs and other computers. Service personnel who monitor and control through HMIs can be called by different names. In industrial process and manufacturing environments, they are called operators or something similar. In boiler houses and central utilities departments they are called stationary engineers.

Different types of automation tools exist:

- ANN - Artificial neural network

- DCS - Distributed Control System
- HMI - Human Machine Interface
- SCADA - Supervisory Control and Data Acquisition
- PLC - Programmable Logic Controller
- PAC - Programmable automation controller
- Instrumentation
- Motion control
- Robotics

Current limits

Many roles for humans in industrial processes presently lie beyond the scope of automation. Human-level pattern recognition, language recognition, and language production ability are well beyond the capabilities of modern mechanical and computer systems. Tasks requiring subjective assessment or synthesis of complex sensory data, such as scents and sounds, as well as high-level tasks such as strategic planning, currently require human expertise. In many cases, the use of humans is more cost-effective than mechanical approaches even where automation of industrial tasks is possible.

Applications of Automation

- **Automated Video surveillance:**

The Defense Advanced Research Projects Agency (DARPA) started the research and development of automated Visual surveillance and Monitoring (VSAM) program 1997-99 and airborne Video Surveillance (AVS) program 1998-2002. Currently there is a major effort underway in the vision community to develop a fully automated tracking surveillance system. Automated video surveillance monitors people and vehicle in real time within a busy environment. Existing automated surveillance systems are based on the environment they are primarily designed to observe, i.e., indoor, outdoor or airborne, the amount of sensors that the automated system can handle and the mobility of sensor, i.e., stationary camera vs. mobile camera. The purpose of a surveillance system is to record properties and trajectories of objects in a given area, generate warnings or notify designated authority in case of occurrence of particular events.

- **Automated Highway Systems:**

As demands for safety and mobility have grown and technological possibilities have multiplied, interest in automation have grown. Seeking to accelerate the development and introduction of fully automated vehicles and highways, The United States Congress authorized more than $650 million over 6 years for intelligent transport systems (ITS) and demonstration projects in the 1991 Intermodal Surface Transportation Efficiency Act (ISTEA). Congress legislated in ISTEA that “The secretary [of transportation] shall develop an automated highway and vehicle prototype from which future fully automated intelligent vehicle-highway systems can be developed. Such development shall include research in human factors to ensure the success of the man-machine relationship. The goal of this program is to have the first fully automated highway roadway or an automated test track in operation by 1997. This system shall accommodate installation of equipment in new and existing motor vehicles." [ISTEA 1991, part B, Section 6054(b)].

Full automation commonly defined as requiring no control or very limited control by the driver; such automation would be accomplished through a combination of sensor, computer, and communications systems in vehicles and along the roadway. Fully automated driving would, in theory, allow closer vehicle spacing and higher speeds, which could enhance traffic capacity in places where additional road building is physically impossible, politically unacceptable, or prohibitively expensive. Automated controls also might enhance road safety by reducing the opportunity for driver error, which causes a large share of motor vehicle crashes. Other potential benefits include improved air quality (as a result of more-efficient traffic flows), increased fuel economy, and spin-off technologies generated during research and development related to automated highway systems.

- **Automated manufacturing:**

Automated manufacturing refers to the application of automation to produce things in the factory way. Most of the advantages of the automation technology has its influence in the manufacture processes.

The main advantage of the automated manufacturing are: higher consistency and quality, reduce the lead times, simplification of production, reduce handling, improve work flow and increase the morale of workers when a good implementation of the automation is made.

- **Home Automation**

Home automation (also called domotics) designates an emerging practice of increased automation of household appliances and features in residential dwellings, particularly through electronic means that allow for things impracticable, overly expensive or simply not possible in recent past decades.

Chapter- 2

Driverless Car

A robotic Volkswagen Passat shown at Stanford University

A **driverless car** is a vehicle equipped with an autopilot system, and capable of driving from one point to another without aid from an operator. Driverless passenger car programs include the 800 million EC EUREKA Prometheus Project on autonomous vehicles, the 2getthere passenger vehicles from the Netherlands, the ARGO research project from Italy, the DARPA Grand Challenge from the USA, and Google driverless car.

History

An early representation of the driverless car was Norman Bel Geddes's Futurama exhibit sponsored by General Motors at the 1933 World's Fair, which depicted electric cars powered by circuits embedded in the roadway and controlled by radio.

The history of autonomous vehicles starts in 1977 with the Tsukuba Mechanical Engineering Lab in Japan. On a dedicated, clearly marked course it achieved speeds of up to 30 km/h (20 miles per hour), by tracking white street markers (special hardware was necessary, since commercial computers were much slower than they are today).

In the 1980s a vision-guided Mercedes-Benz robot van, designed by Ernst Dickmanns and his team at the Bundeswehr University of Munich in Munich, Germany, achieved 100 km/h on streets without traffic. Subsequently, the European Commission began funding the 800 million Euro EUREKA Prometheus Project on autonomous vehicles (1987–1995).

Also in the 1980s the DARPA-funded Autonomous Land Vehicle (ALV) in the United States achieved the first road-following demonstration that used laser radar (Environmental Research Institute of Michigan), computer vision (Carnegie Mellon University and SRI), and autonomous robotic control (Carnegie Mellon and Martin Marietta) to control a driverless vehicle up to 30 km/h. In 1987, HRL Laboratories (formerly Hughes Research Labs) demonstrated the first off-road map and sensor-based autonomous navigation on the ALV. The vehicle travelled over 600m at 3 km/h on complex terrain with steep slopes, ravines, large rocks, and vegetation.

In 1994, the twin robot vehicles VaMP and Vita-2 of Daimler-Benz and Ernst Dickmanns of UniBwM drove more than one thousand kilometers on a Paris three-lane highway in standard heavy traffic at speeds up to 130 km/h, albeit semi-autonomously with human interventions. They demonstrated autonomous driving in free lanes, convoy driving, and lane changes left and right with autonomous passing of other cars.

In 1995, Dickmanns´ re-engineered autonomous S-Class Mercedes-Benz took a 1600 km trip from Munich in Bavaria to Copenhagen in Denmark and back, using saccadic computer vision and transputers to react in real time. The robot achieved speeds exceeding 175 km/h on the German Autobahn, with a mean time between human interventions of 9 km, or 95% autonomous driving. Again it drove in traffic, executing manoeuvres to pass other cars. Despite being a research system without emphasis on long distance reliability, it drove up to 158 km without human intervention.

In 1995, the Carnegie Mellon University Navlab project achieved 98.2% autonomous driving on a 5000 km (3000-mile) "No hands across America" trip. This car, however, was semi-autonomous by nature: it used neural networks to control the steering wheel, but throttle and brakes were human-controlled.

From 1996–2001, Alberto Broggi of the University of Parma launched the ARGO Project, which worked on enabling a modified Lancia Thema to follow the normal (painted) lane marks in an unmodified highway. The culmination of the project was a journey of 2,000 km over six days on the motorways of northern Italy dubbed MilleMiglia in Automatico, with an average speed of 90 km/h. 94% of the time the car was in fully automatic mode, with the longest automatic stretch being 54 km. The vehicle had only two black-and-white low-cost video cameras on board, and used stereoscopic vision algorithms to understand its environment, as opposed to the "laser, radar - whatever you need" approach taken by other efforts in the field.

Three US Government funded military efforts known as Demo I (US Army), Demo II (DARPA), and Demo III (US Army), are currently underway. Demo III (2001) demonstrated the ability of unmanned ground vehicles to navigate miles of difficult off-road terrain, avoiding obstacles such as rocks and trees. James Albus at NIST provided the Real-Time Control System which is a hierarchical control system. Not only were individual vehicles controlled (e.g. throttle, steering, and brake), but groups of vehicles had their movements automatically coordinated in response to high level goals.

In 2002, the DARPA Grand Challenge competitions were announced. The 2004 and 2005 DARPA competitions allowed international teams to compete in fully autonomous vehicle races over rough unpaved terrain and in a non-populated suburban setting. The 2007 DARPA challenge, the DARPA urban challenge, involved autonomous cars driving in an urban setting.

In 2008, General Motors stated that they will begin testing driverless cars by 2015, and they could be on the road by 2018 .

In 2010 VisLab ran VIAC, the VisLab Intercontinental Autonomous Challenge, a 13,000 km test run of autonomous vehicles. The four driverless electric vans successfully ended the drive from Italy to China via the arriving at the Shanghai Expo on 28 October.

Recent projects

The work done so far varies significantly in its ambition and its demands in terms of modification of the infrastructure. Broadly, there are three approaches:

- Fully autonomous vehicles
- Various enhancements to the infrastructure (either an entire area, or specific lanes) to create a self-driving closed system.
- "assistance" systems that incrementally remove requirements from the human driver (e.g. improvements to cruise control)

An important concept that cuts across several of the efforts is vehicle **platoons**. In order to better utilize road-space, vehicles are assembled into ad-hoc train-like "platoons", where the driver (either human or automatic) of the first vehicle makes all decisions for the entire platoon. All other vehicles simply follow the lead of the first vehicle.

Fully autonomous

Fully autonomous driving requires a car to drive itself to a pre-set target using unmodified infrastructure. The final goal of safe door-to-door transportation in arbitrary environments is not yet reached though.

Vehicles for paved roads

- Google driverless car, with a test fleet of autonomous vehicles that by October 2010 have driven 140,000 miles (230,000 km) without any incidents.
- The 800 million Euro EUREKA Prometheus Project on autonomous vehicles (1987–1995). Among its culmination points were the twin robot vehicles VITA-2 and VaMP of Daimler-Benz and Ernst Dickmanns, driving long distances in heavy traffic.
- The VIAC Challenge, in which 4 vehicles drove from Italy to China on a 13,000 kilometres (8,100 mi) trip with only limited occasions intervene by human, such as in the Moscow traffic jams and when passing toll stations. This is the longest-ever trip by an unmanned vehicle.
- The third competition of the DARPA Grand Challenge held in November 2007. 53 teams qualified initially, but after a series of qualifying rounds, only eleven teams entered the final race. Of these, six teams completed navigating through the non-populated urban environment, and the Carnegie Mellon University team won the $2 million prize.
- The ARGO vehicle is the predecessor of the BRAiVE vehicle, both from the University of Parma's VisLab. Argo was developed in 1996 and demonstrated to the world in 1998; BRAiVE was developed in 2008 and firstly demonstrated in 2009 at the IEEE IV conference in Xi'an, China.
- Stanford Racing Team's Junior car is an autonomous driverless car for paved roads. It is intended for civilian use.
- The Volkswagen Golf GTI 53+1 is a modified Volkswagen Golf GTI capable of autonomous driving. The Golf GTI 53+1 features a implemented system that can be integrated into any car. This system is based around the MicroAutoBox from dSpace. This, as it was intended to test VW hardware without a human driver (for consistent test results).
- The Audi TTS Pikes Peak is a modified Audi TTS, working entirely on GPS, and thus without additional sensors. The car was designed by Burkhard Huhnke of Volkswagen Research.
- Stadtpilot, Technical University Braunschweig
- AutoNOMOS - part of the Artificial Intelligence Group of the Freie Universität Berlin

Free-ranging vehicles

There are three clusters of activity relating to free-ranging off-road cars. Some of these projects are military-oriented.

- US military DARPA Grand Challenge

  The US Department of Defense announced on the July 30, 2002 a "Grand Challenge", for US-based teams to produce a vehicle that could autonomously navigate and reach a target in the desert of the south western USA.
  In March 2004, the first competition was held, for a prize-money of $1 million. Not one of the 25 entrants completed the course. However, in the second competition held in October 2005 five different teams completed the 135-mile (217 km) course, and the Stanford University team won the $2 million prize. November 3rd, 2007, the third competition was held and $3.5 million dollar in cash prizes, trophies and medals were awarded. Six driverless vehicles were able to complete the 55 miles (89 km) of urban traffic in the 2007 DARPA Urban Challenge rally style race. 1st Place - Tartan Racing, Pittsburgh, PA; 2nd Place - Stanford Racing Team, Stanford, CA; 3rd Place - Victor Tango, Blacksburg, VA.

- European Land-Robot Trial (ELROB)

  The German Department of Defense held an exhibition trade show (ELROB) for demonstrating automated vehicles in May 2006. The event included various military automated and remotely-operated robots, for various military uses. Some of the systems on display could be ordered and implemented immediately. In August 2007 a civilian version of the event was held in Switzerland.
  The Smart team from Switzerland presented "a Vehicle for Autonomous Navigation and Mapping in Outdoor Environments".

- The Israeli Military-Industrial Complex

  As a followup from its success with Unmanned Combat Air Vehicles, and following the construction of the Israeli West Bank barrier there has been significant interest in developing a fully automated border-patrol vehicle. Two projects, by Elbit Systems and Israel Aircraft Industries are both based on the locally-produced Armored "Tomcar" and have the specific purpose of patrolling barrier fences against intrusions.
  The "SciAutonics II" team in the 2004 DARPA Challenge used Elbit's version of the Tomcar.

Pre-built infrastructure

The following projects were conceived as practical attempts to use available technology in an incremental manner to solve specific problems, like transport within a defined campus area, or driving along a stretch of motorway. The technologies are proven, and the main barrier to widespread implementation is the cost of deploying the infrastructure. Such systems already function in many airports, on railroads, and in some European towns.

Dual mode transit - monorail

There is a family of projects, all currently still at the experimental stage, that would combine the flexibility of a private automobile with the benefits of a monorail system. The idea is that privately-owned cars would be built with the ability to dock themselves onto a public monorail system, where they become part of a centrally managed, fully computerized transport system—more akin to a driverless train system (as already found in airports) than to a driverless car. This idea is also known as Dual mode transit.

Groups working on this concept are:

- RUF (Denmark)
- BiWay (UK)
- ATN (New Zealand)
- TriTrack (Texas, United States)

Automated highway systems

Automated highway systems (**AHS**) are an effort to construct special lanes on existing highways that would be equipped with magnets or other infrastructure to allow vehicles to stay in the center of the lane, while communicating with other vehicles (and with a central system) to avoid collision and manage traffic. Like the dual-mode monorail, the idea is that cars remain private and independent, and just use the AHS system as a quick way to move along designated routes. AHS allows specially equipped cars to join the system using special 'acceleration lanes' and to leave through 'deceleration lanes'. When leaving the system each car verifies that its driver is ready to take control of the vehicle, and if that is not the case, the system parks the car safely in a predesignated area.

Some implementations use radar to avoid collisions and coordinate speed.

One example that uses this implementation is the AHS demo of 1997 near San Diego, sponsored by the US government, in coordination with the State of California and Carnegie Mellon University. The test site is a 12-kilometer, high-occupancy-vehicle (HOV) segment of Interstate 15, 16 kilometers north of downtown San Diego. The event generated much press coverage.

This concerted effort by the US government seems to have been pretty much abandoned because of social and political forces, above all else the desire to create a less futuristic and more marketable solution.

As of 2007, a three-year project is underway to allow robot controlled vehicles, including buses and trucks, to use a special lane along 20 Interstate 805. The intention is to allow the vehicles to travel at shorter following distances and thereby allow more vehicles to use the lanes. The vehicles will still have drivers since they need to enter and exit the special lanes. The system is being designed by Swoop Technology, based in San Diego county.

Free-ranging on grid

Frog Navigation Systems (the Netherlands) applies the FROG (**free-ranging on grid**) technology. The technology consists of a combination of autonomous vehicles and a supervisory central system. The company's purpose-built electric vehicles locate themselves using odometry readings, recalibrating themselves occasionally using a "maze" of magnets embedded in the environment, and GPS. The cars avoid collisions with obstacles located in the environment using laser (long range) and ultra-sonic (short-range) sensors.

The vehicles are completely autonomous and plan their own routes from A to B. The supervisory system merely administers the operations and directs traffic where required. The system has been applied both indoors and outdoors, and in environments where 100+ automated vehicles are operational (container port). At this time the system is not suited yet for running the sheer number of vehicles encountered in urban settings. The company also has no intention of developing such technology at this time.

The FROG system is deployed for industrial purposes in factory sites, and is marketed as a pilot public transport system in the city of Capelle aan den IJssel by its subsidiary 2getthere. This system experienced an accident that proved to be caused by a Human error.

Frog Navigation Systems is one of few fully commercial companies in this field.

Driver-assistance

Though these products and projects do not aim explicitly to create a fully autonomous car, they are seen as incremental stepping-stones in that direction. Many of the technologies detailed below will probably serve as components of any future driverless car — meanwhile they are being marketed as gadgets that assist human drivers in one way or another. This approach is slowly trickling into standard cars (e.g. improvements to cruise control).

Driver-assistance mechanisms are of several distinct types, sensorial-informative, actuation-corrective, and systemic.

Sensorial-informative

These systems warn or inform the driver about events that may have passed unnoticed, such as

- Lane Departure Warning System (LDWS), for example from Iteris or Mobileye N.V.
- Rear-view alarm, to detect obstacles behind.
- Visibility aids for the driver, to cover blind spots and enhanced vision systems such as radar, wireless vehicle safety communications and night vision.

- Infrastructure-based, driver warning/information-giving systems, such as those developed by the Japanese government

Actuation-corrective

These systems modify the driver's instructions so as to execute them in a more effective way, for example the most widely deployed system of this type is ABS; conversely power steering is not a control mechanism, but just a convenience - it is not involved in decision making.

- Anti-lock braking system (ABS) (also Emergency Braking Assistance (EBA), often coupled with Electronic brake force distribution (EBD), which prevents the brakes from locking and losing traction while braking. This shortens stopping distances in most cases and, more importantly, allows the driver to steer the vehicle while braking.
- Traction control system (TCS) actuates brakes or reduces throttle to restore traction if driven wheels begin to spin.
- Four wheel drive (AWD) with a centre differential. Distributing power to all four wheels lessens the chances of wheel spin. It also suffers less from oversteer and understeer.
- Electronic Stability Control (ESC) (also known for Mercedes-Benz proprietary Electronic Stability Program (ESP), Acceleration Slip Regulation (ASR) and Electronic differential lock (EDL)). Uses various sensors to intervene when the car senses a possible loss of control. The car's control unit can reduce power from the engine and even apply the brakes on individual wheels to prevent the car from understeering or oversteering.
- Dynamic steering response (DSR) corrects the rate of power steering system to adapt it to vehicle's speed and road conditions.

A review of the overall "feel" to actuation-correction in a Jaguar XK convertible.

Driver-assistance preview from Popular Science (dated 2004).

Note: The electronic differential lock (EDL) employed by Volkswagen is not - as the name suggests - a differential lock at all. Sensors monitor wheel speeds, and if one is rotating substantially faster than the other (i.e. slipping) the EDL system momentarily brakes it. This effectively transfers all the power to the other wheel.

Systemic

- Automatic parking: e.g. technology from Ford or Toyota selling for $700, with a 70% take-up rate. The Lexus LS can park itself (parallel/reverse) via the 'Advanced Parking Guidance System' – though only controlling the steering.
- Follow another car on a motorway ("Enhanced" or "adaptive" cruise control), like The Ford or Vauxhall(GM).
- Nissan's "Distance Control assist"

- Dead Man's Switch; there is a move to introduce deadman's braking into automotive application, primarily heavy vehicles, and there may also be a need to add penalty switches to cruise controls.

Existing and missing technologies

In order to drive a car, a system would need to:

1. Understand its immediate environment (Sensors)
2. Know where it is and where it wants to go (Navigation)
3. Find its way in the traffic (Motion planning)
4. Operate the mechanics of the vehicle (Actuation)

Arguably, 2½ of these problems are already solved: Navigation and Actuation completely, and Sensors partially, but improving fast. The main unsolved part is the motion planning.

Sensors

Sensors employed in driverless cars vary from the minimalist ARGO project's monochrome stereoscopy to Mobileye's inter-modal (video, infra-red, laser, radar) approach. The minimalist approach imitates the human situation most closely, while the multi-modal approach is "greedy" in the sense that it seeks to obtain as much information as is possible by current technology, even at the occasional cost of one car's detection system interfering with another's.

Mobileye N.V. is a technology company that focuses on the development of vision-based Advanced Driver Assistance Systems (ADAS) providing warnings for collision prevention and mitigation. Mobileye offers a wide range of driver safety solutions combining artificial vision image processing, multiple technological applications and information technology. Mobileye's vehicle detection systems, are currently only used for driver assistance, but are eminently suitable for a full-fledged driverless car. This video demonstrates the capabilities of the system: all pedestrians, cars, motorbikes etc. are clearly displayed in video, with a frame around them and the distance between "our" car and the object observed. The system also detects the objects' motion (direction and speed) and can so calculate relative speeds, and predict collisions.

- Japanese infra-red article
- some things from the DARPA challenge....
- Road-sign recognition

Navigation

The ability to plot a route from where the vehicle is to where the user wants to be has been available for several years. These systems, based on the US military's Global Positioning System are now available as standard car fittings, and use satellite transmissions to ascertain the current location, and an on-board street database to derive a route to the target. The more sophisticated systems also receive radio updates on road blockages, and adapt accordingly. There are also sensors that greatly affect the whole nature of it.

Motion planning

- PMP + SLAMMOT YouTube

This is current research problem.

Control of vehicle

As automotive technology matures, more and more functions of the underlying engine, gearbox etc. are no longer directly controlled by the driver by mechanical means, but rather via a computer, which receives instructions from the driver as inputs and delivers the desired effect by means of electronic throttle control, and other drive-by-wire elements. Therefore, the technology for a computer to control all aspects of a vehicle is well understood.

Work done in simulation

While developing control systems for real cars is very costly in terms of both time and money, much work can be done in simulations of various complexity. Systems developed using simpler simulators can gradually be transferred to more complex simulators, and in the end to real vehicles.

Social impact

Driverless cars may yield advantages of increasing roadway capacity by reducing the distances between cars, reduce congestion by efficiently controlling the flow of traffic, and increase safety by eliminating driver error.

According to urban designer and futurist Michael E. Arth, driverless electric vehicles—in conjunction with the increased use of virtual reality for work, travel, and pleasure—could reduce the world's vehicles (estimated to be 800,000,000) to a fraction of that number within a few decades. Arth claims that this would be possible if almost all private cars requiring drivers, which are not in use and parked 90% of the time, would be traded for public self-driving taxis that would be in near constant use. This would also allow for getting the appropriate vehicle for the particular need—a bus could come for a group of people, a limousine could come for a special night out, and a Segway could come for a short trip down the street for one person. Children could be chauffeured in supervised safety, DUIs would no longer exist, and 41,000 lives could be saved each year in the U.S. alone.

Key players

International

The European Union has a multi-billion Euro programme to support Research and Development by ad-hoc consortia from the various member countries, called Framework Programmes for Research and Technological Development. Several of these projects pertain to the subject of driverless cars, e.g.:

- INRIA's La Route Automatisée project gathered much useful data about the actual and possible deployments of Driverless Cars for **public transport**. The main system discussed is based on FROG.

Many of the EU-sponsored projects are coordinated by a group called Ertico.

There are several national associations around the world that are active in research in the field of intelligent transportation systems, a term that seems to encompass anything which applies technology to the improvement of transport. In recent years there has been a trend in this field to move efforts away from the more visionary projects, such as driverless cars, to the more short-term, such as public transport and traffic management. Many of these organizations are government sponsored, and they all cooperate at some level or another. Some of the countries involved are: USA, IEEE ITS Society, Australia, South Korea, Taiwan, India--(specifically Intelligent vehicles), and Japan

Chapter- 3

Robotic Surgery

Robotic surgery, computer-assisted surgery, and robot-assisted surgery are terms for various technological developments that currently are developed to support a range of surgical procedures.

Robot-assisted surgery was developed to overcome limitations of minimally invasive surgery, Instead of directly moving the instruments the surgeon uses a computer console to manipulate the instruments attached to multiple robot arms. The computer translates the surgeon's movements, which are then carried out on the patient by the robot. Other features of the robotic system include, for example, an integrated tremor filter and the ability for scaling of movements (changing of the ratio between the extent of movements at the master console to the internal movements of the instruments attached to the robot). The console is located in the same operating room as the patient, but is physically separated from the operative workspace. Since the surgeon does not need to be in the immediate location of the patient while the operation is being performed, it can be possible for specialists to perform remote surgery on patients. Robots can perform surgery without a human surgeon.

History

The world's first surgical robot was the "Arthrobot", which was developed and used for the first time in Vancouver, BC, Canada in 1983. The robot was developed by a team led by Dr. James McEwen and Geof Auchinlek, in collaboration with orthopaedic surgeon, Dr. Brian Day. National geographic produced a movie on robotics which featured the Arthrobot. In related projects at that time, other medical robots were developed, including a robotic arm that performed eye surgery, and another that acted as an operating assistant, and handed the surgeon instruments in response to voice commands.

1985 a robot, the PUMA 560, was used to place a needle for a brain biopsy using CT guidance. In 1988, the PROBOT, developed at Imperial College London, was used to perform prostatic surgery. The ROBODOC from Integrated Surgical Systems was introduced in 1992 to mill out precise fittings in the femur for hip replacement. Further development of robotic systems was carried out by Intuitive Surgical with the introduction of the da Vinci Surgical System and Computer Motion with the *AESOP* and the ZEUS robotic surgical system. (Intuitive Surgical bought Computer Motion in 2003; ZEUS is no longer being actively marketed.)

The da Vinci Surgical System comprises three components: a surgeon’s console, a patient-side robotic cart with 4 arms manipulated by the surgeon (one to control the camera and three to manipulate instruments), and a high-definition 3D vision system. Articulating surgical instruments are mounted on the robotic arms which are introduced into the body through cannulas. The device senses the surgeon’s hand movements and translates them electronically into scaled-down micro-movements to manipulate the tiny proprietary instruments. It also detects and filters out any tremors in the surgeon's hand movements, so that they are not duplicated robotically. The camera used in the system provides a true stereoscopic picture transmitted to a surgeon's console. The da Vinci System is FDA cleared for a variety of surgical procedures including surgery for prostate cancer, hysterectomy and mitral valve repair, and is used in more than 800 hospitals in the Americas and Europe. The da Vinci System was used in 48,000 procedures in 2006 and sells for about $1.2 million. The new da Vinci HD SI released in April, 2009 currently sells for $1.75 million. The first robotic surgery took place at The Ohio State University Medical Center in Columbus, Ohio under the direction of Dr. Robert E. Michler, Professor and Chief, Cardiothoracic Surgery.<McConnell PI, Schneeberger EW, Michler RE. History and development of robotic cardiac surgery. Problems in General Surgery 2003;20:62-72.>

In September 2010, the Eindhoven University of Technology announced the development of the Sofie surgical system, the first surgical robot to employ force feedback.

- In 1997 a reconnection of the fallopian tubes operation was performed successfully in Cleveland using ZEUS.
- In May 1998, Dr. Friedrich-Wilhelm Mohr using the *Da Vinci surgical robot* performed the first robotically assisted heart bypass at the Leipzig Heart Centre in Germany.
- On 2 September 1999, Dr. Randall Wolf and Dr. Robert Michler performed the first robotically assisted heart bypass in the USA at The Ohio State University.
- In October 1999 the world's first surgical robotics *beating heart* coronary artery bypass graft (CABG) was performed in Canada by Dr. Douglas Boyd and Dr. Reiza Rayman using the ZEUS surgical robot.
- In 2001, Prof. Marescaux, while in New York, used the "Zeus" robot to remotely perform gall bladder surgery on a patient who was in Strasbourg, France.
- In September 2001, Dr. Michel Gagner, while in New York, used the Zeus robotic system to remotely perform a cholecystectomy on a woman who was in Strasbourg, France.
- In May 2006 the first AI doctor-conducted unassisted robotic surgery on a 34 year old male to correct heart arythmia. The results were rated as better than an above-average human surgeon. The machine had a database of 10,000 similar

operations, and so, in the words of its designers, was "more than qualified to operate on any patient." The designers believe that *robots can replace half of all surgeons within 15 years*.

- In February 2008, Dr. Mohan S. Gundeti of the University of Chicago Comer Children's Hospital performed the first robotic pediatric neurogenic bladder reconstruction. The operation was performed on a 10-year-old girl.
- In January 2009, Dr. Todd Tillmanns reported the results of the largest multi-institutional study on the use of the da-Vinci robotic surgical system in gynecologic oncology and included learning curves for current and new users as a method to assess their acquisition of skills using the device.
- In January 2009, the first all-robotic-assisted kidney transplant was performed at Saint Barnabas Medical Center in Livingston, New Jersey by Dr. Stuart Geffner. The same team performed eight more fully robotic-assisted kidney transplants over the next six months.

Advantages and disadvantages

Major advances aided by surgical robots have been remote surgery, minimally invasive surgery and unmanned surgery. Some major advantages of robotic surgery are precision, miniaturization, smaller incisions, decreased blood loss, less pain, and quicker healing time. Further advantages are articulation beyond normal manipulation and three-dimensional magnification, resulting in improved ergonomics. Robotic techniques are also associated with reduced duration of hospital stays, blood loss, transfusions, and use of pain medication.

With a the cost of the robot at $1,200,000 dollars and disposable supply costs of $1,500 per procedure, the cost of the procedure is higher. Additional surgical training is needed to operate the system. Patient surveys indicate they chose the procedure based on expectations of decreased morbidity, improved outcomes, reduced blood loss and less pain. Higher expectations may explain higher rates of dissatisfaction and regret.

The main advantage of this technique is that the incisions are very small and, consequently, patient recovery is quick. In traditional open-heart surgery, the surgeon makes a ten to twelve-inch incision, then accesses the heart by splitting the sternum (breast bone) and spreading open the rib cage. The patient is then placed on a heart-lung machine and the heart is stopped for the length of the surgery. Not only is this a way for bacteria that can cause infections to access the patient's body, it also leads to a painful wound, which takes time to heal.

Because patient recovery after robot-assisted heart surgery is quicker, the hospital stay is shorter. On average patients leave the hospital two to five days earlier than patients who have undergone traditional open-heart surgery and return to work and normal activity 50% more quickly. Reduced recovery times are not only better for the patient, they also reduce the number of staff needed during surgery, nursing care required after surgery, and, therefore, the overall cost of hospital stays.

Compared with other minimally invasive surgery approaches, robot-assisted surgery gives the surgeon better control over the surgical instruments and a better view of the surgical site. In addition, surgeons no longer have to stand throughout the surgery and do not tire as quickly. Naturally occurring hand tremors are filtered out by the robot's computer software. Finally, the surgical robot can continuously be used by rotating surgery teams (Gerhardus 2003). Gerhardus D (2003). Robot-assisted surgery: the future is here. Journal of Healthcare Management, Jul/Aug, 242-251. While the use of robotic surgery has become a item in the advertisement of medical services, critics point out that studies that indicate that long-term results are superior to those after laparoscopic surgery are lacking. The robotic system does not come cheap and has a learning curve. Data are absent to show that the increased costs can be justified. In the medical literature, very experienced surgeons tend to publish their results, these, however, may not be representative of surgeons with lesser experience.

The cost of robotic surgical systems lies between $750.000 and $1.2 million (as of 2005). Numerous financial feasibility studies have been done to determine whether it is really worth a hospital's while to purchase such a system and opinions differ dramatically. Surgeons report that, although the manufacturers of the systems provide training on this new technology, the learning phase is intensive and surgeons must operate on twelve to eighteen patients before they feel comfortable with the system. During the training phase, minimally invasive operations can take up to twice as long as traditional surgery, which ties up operating room and surgical staff time and keeps patients under anesthesia longer.

Applications

General surgery

In 2007, the. University of Illinois at Chicago medical team, lead by Prof. Pier Cristoforo Giulianotti, performed the world's first ever robotic pancreatectomy and also the Midwests fully robotic Whipple surgery. In April 2008, the same team of surgeons performed the world's first fully minimally invasive liver resection for living donor transplantation, removing 60% of the patient's liver, yet allowing him to leave the hospital just a couple of days after the procedure, in very good condition. Furthermore the patient can also leave with less pain than a usual surgery due to the four puncture holes and not a scar by a surgeon.

Cardiothoracic surgery

Robot-assisted MIDCAB and Endoscopic coronary artery bypass (TECAB) surgeries are being performed with the da Vinci system. Mitral valve repairs and replacements have been performed. East Carolina University, Greenville (Dr W. Randolph Chitwood), Saint Joseph's Hospital, Atlanta (Dr Douglas A. Murphy), and Good Samaritan Hospital, Cincinnati (Dr J. Michael Smith) have popularized this procedure and proved its durability with multiple publications. Since the first robotic cardiac procedure performed in the USA in 1999, The Ohio State University, Columbus (Dr. Robert E. Michler, Dr. Juan Crestanello, Dr. Paul Vesco) has performed CABG, mitral valve, esophagectomy, lung resection, tumor resections, among other robotic assisted procedures and serves as a training site for other surgeons. In 2002, surgeons at the Cleveland Clinic in Florida (Dr. Douglas Boyd and Kenneth Stahl) reported and published their preliminary experience with minimally invasive "hybrid" procedures. These procedures combined robotic revascularization and coronary stenting and further expanded the role of robots in coronary bypass to patients with disease in multiple vessels.

Cardiology and electrophysiology

The Stereotaxis Magnetic Navigation System (MNS) has been developed to increase precision and safety in ablation procedures for arrhythmias and atrial fibrillation while reducing radiation exposure for the patient and physician, and the system utilizes two magnets to remotely steerable catheters. The system allows for automated 3-D mapping of the heart and vasculature, and MNS has also been used in interventional cardiology for guiding stents and leads in PCI and CTO procedures, proven to reduce contrast usage and access tortuous anatomy unreachable by manual navigation. Dr. Andrea Natale has referred to the new Stereotaxis procedures with the magnetic irrigated catheters as "revolutionary."

The Hansen Medical Sensei robotic catheter system uses a remotely operated system of pulleys to navigate a steerable sheath for catheter guidance. It allows precise and more forceful positioning of catheters used for 3-D mapping of the heart and vasculature. The system provides doctors with estimated force feedback information and feasible manipulation within the left atrium of the heart. The Sensei has been associated with mixed acute success rates compared to manual, commensurate with higher procedural complications, longer procedure times but lower fluoroscopy dosage to the patient.

It was estimated that 70 to 90 hospitals in the United States now use minimally invasive surgical robots for heart surgery, and this number is expected to double by mid-2006 (Alt and Worrell 2004). At present, three types of heart surgery are being performed on a routine basis using robotic surgery systems (Kypson and Chitwood 2004). These three surgery types are:

Atrial septal defect repair — the repair of a hole between the two upper chambers of the heart, Mitral valve repair — the repair of the valve that prevents blood from regurgitating back into the upper heart chambers during contractions of the heart, Coronary artery bypass — rerouting of blood supply by bypassing blocked arteries that provide blood to the heart.

As surgeons' experience and robotic technology develop, it is expected that robot-assisted procedures will be applied to additional types of heart surgery.

Alt SJ & Worrell B (2004). More surgeons do minimally invasive heart surgery. Health Care Strategic Management, Apr, 1 & 11-19.

Kypson AP & Chitwood WR Jr. (2004). Robotic applications in cardiac surgery. International Journal of Advanced Robotic Systems, 1(2), 87-92.

Gastrointestinal surgery

Multiple types of procedures have been performed with either the *Zeus* or *da Vinci robot* systems, including bariatric surgery.

Gynecology

Robotic surgery in gynecology is one of the fastest growing fields of robotic surgery. This includes the use of the da Vinci surgical system in benign gynecology and gynecologic oncology. Robotic surgery can be used to treat fibroids, abnormal periods, endometriosis, ovarian tumors, pelvic prolapse, and female cancers. Using the robotic system, gynecologists can perform hysterectomies, myomectomies, and lymph node biopsies. The need for large abdominal incisions is virtually eliminated.

Robot assisted hysterectomies and cancer staging are being performed using da Vinci robotic system. The University of Tennessee, Memphis (Dr. Todd Tillmanns, Dr. Saurabh Kumar), Northwestern University (Dr. Patrick Lowe), Aurora Health Center (Dr. Scott Kamelle), West Virginia University (Dr. Jay Bringman) and The University of Tennessee, Chattanooga (Dr. Donald Chamberlain) have extensively studied the use of robotic surgery and found it to improve morbidity and mortality of patients with gynecologic cancers. They have also for the first time reported robotic surgery learning curves for current and new users as a method to assess acquisition of their skills using the device.

Neurosurgery

Several systems for stereotactic intervention are currently on the market. MD Robotic's NeuroArm is the world's first MRI-compatible surgical robot.

Orthopedics

The ROBODOC system was released in 1992 by Integrated Surgical Systems, Inc. which merged into CUREXO Technology Corporation. Also, The Acrobot Company Ltd. sells the "Acrobot Sculptor", a robot that constrains a bone cutting tool to a pre-defined volume. Another example is the CASPAR robot produced by U.R.S.-Ortho GmbH & Co. KG, which is used for total hip replacement, total knee replacement and anterior cruciate ligament reconstruction.

Pediatrics

Surgical robotics has been used in many types of pediatric surgical procedures including: tracheoesophageal fistula repair, cholecystectomy, nissen fundoplication, morgagni's hernia repair, kasai portoenterostomy, congenital diaphragmatic hernia repair, and others.

On January 17, 2002, surgeons at Children's Hospital of Michigan in Detroit performed the nation's first advanced computer-assisted robot-enhanced surgical procedure at a children's hospital.

The Center for Robotic Surgery at Children's Hospital Boston provides a high level of expertise in pediatric robotic surgery. Specially-trained surgeons use a high-tech robot to perform complex and delicate operations through very small surgical openings. The results are less pain, faster recoveries, shorter hospital stays, smaller scars, and happier patients and families.

In 2001, Children's Hospital Boston was the first pediatric hospital to acquire a surgical robot. Today, surgeons use the technology for many procedures and perform more pediatric robotic surgeries than any other hospital in the world. Children's Hospital physicians have developed a number of new applications to expand the use of the robot, and train surgeons from around the world on its use.

Radiosurgery

The CyberKnife Robotic Radiosurgery System uses image-guidance and computer controlled robotics to treat tumors throughout the body by delivering multiple beams of high-energy radiation to the tumor from virtually any direction.

Urology

Removing the prostate gland for cancer, repair obstructed kidneys, repair bladder abnormalities and remove diseased kidneys. New minimally invasive robotic devices using steerable flexible needles are currently being developed for use in prostate brachytherapy. A few leading urologists in the field of robotic urological surgery are Drs. David Samadi, Ashutosh Tewari, Mani Menon, Peter Schlegel, Mehmood Akhtar, Douglas Scherr, Mohamad W. Salkini, Steven Sukin, and Vipul Patel.

In 2000, the first robot-assisted laparoscopic radical prostatectomy was performed.

Miniature robotics

As scientists seek to improve the versatility and utility of robotics in surgery, some are attempting to miniaturize the robots. For example, the University of Nebraska Medical Center has led a multi-campus effort to provide collaborative research on mini-robotics among surgeons, engineers and computer scientists. There may also be a day and age where nanorobots may be inserted into peoples bloodstreams to act as General Practitioners, or **GPs**; Analysing the problem and sending the information back to the hospital. This could one day remove the need of GPs.

Chapter- 4

Computer Animation (Application of Motion Planning)

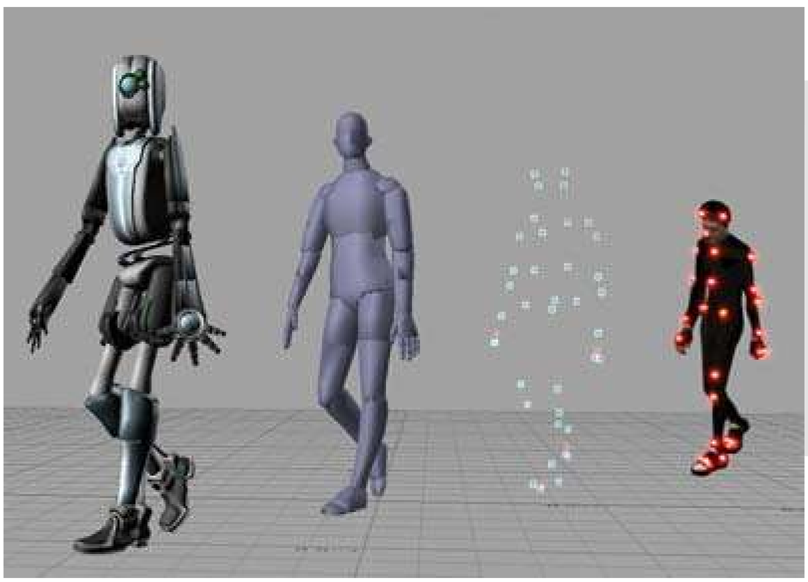

An example of computer animation which is produced in the "motion capture" technique

Computer animation is the process used for generating animated images by using computer graphics. The more general term computer generated imagery encompasses both static scenes and dynamic images, while *computer animation* only refers to moving images produced by exploiting the persistence of vision to make a series of images look animated. Given that images last for about one twenty-fifth of a second on the retina fast image replacement creates the illusion of movement.

Modern computer animation usually uses 3D computer graphics, although 2D computer graphics are still used for stylistic, low bandwidth, and faster real-time renderings. Sometimes the target of the animation is the computer itself, but sometimes the target is another medium, such as film.

Computer animation is essentially a digital successor to the art of stop motion animation of 3D models and frame-by-frame animation of 2D illustrations. Computer generated animations are more controllable than other more physically based processes, such as constructing miniatures for effects shots or hiring extras for crowd scenes, and because it allows the creation of images that would not be feasible using any other technology. It can also allow a single graphic artist to produce such content without the use of actors, expensive set pieces, or props.

To create the illusion of movement, an image is displayed on the computer screen and repeatedly replaced by a new image that is similar to the previous image, but advanced slightly in the time domain (usually at a rate of 24 or 30 frames/second). This technique is identical to how the illusion of movement is achieved with television and motion pictures.

For 3D animations, objects (models) are built on the computer monitor (modeled) and 3D figures are rigged with a virtual skeleton. For 2D figure animations, separate objects (illustrations) and separate transparent layers are used, with or without a virtual skeleton. Then the limbs, eyes, mouth, clothes, etc. of the figure are moved by the animator on key frames. The differences in appearance between key frames are automatically calculated by the computer in a process known as tweening or morphing. Finally, the animation is rendered.

For 3D animations, all frames must be rendered after modeling is complete. For 2D vector animations, the rendering process is the key frame illustration process, while tweened frames are rendered as needed. For pre-recorded presentations, the rendered frames are transferred to a different format or medium such as film or digital video. The frames may also be rendered in real time as they are presented to the end-user audience. Low bandwidth animations transmitted via the internet (e.g. 2D Flash, X3D) often use software on the end-users computer to render in real time as an alternative to streaming or pre-loaded high bandwidth animations.

Explanation

To trick the eye and brain into thinking they are seeing a smoothly moving object, the pictures should be drawn at around 12 frames per second (frame/s) or faster (a frame is one complete image). With rates above 70 frames/s no improvement in realism or smoothness is perceivable due to the way the eye and brain process images. At rates below 12 frame/s most people can detect jerkiness associated with the drawing of new images which detracts from the illusion of realistic movement. Conventional hand-drawn cartoon animation often uses 15 frames/s in order to save on the number of drawings needed, but this is usually accepted because of the stylized nature of cartoons. Because it produces more realistic imagery computer animation demands higher frame rates to reinforce this realism.

The reason no jerkiness is seen at higher speeds is due to “persistence of vision.” From moment to moment, the eye and brain working together actually store whatever one looks at for a fraction of a second, and automatically "smooth out" minor jumps. Movie film seen in theaters in the United States runs at 24 frames per second, which is sufficient to create this illusion of continuous movement.

History

CGI was first used in movies in 1973's *Westworld,* a science-fiction film about a society in which robots live and work among humans, though the first use of 3D Wireframe imagery was in its sequel, *Futureworld* (1976), which featured a computer-generated hand and face created by then University of Utah graduate students Edwin Catmull and Fred Parke. The third movie to use this technology was *Star Wars* (1977) for the scenes with the wireframe Death Star plans and the targeting computers in the X-wings and the *Millennium Falcon*. *The Black Hole* (1979) used raster wire-frame model rendering to depict a black hole. The science fiction-horror film *Alien* of that same year also used a raster wire-frame model, in this case to render the image of navigation monitors in the sequence where a spaceship follows a beacon to a land on an unfamiliar planet.

In 1978, graduate students at the New York Institute of Technology Computer Graphics Lab began work on what would have been the first full-length CGI film, *The Works*, and a trailer for it was shown at SIGGRAPH 1982, but the film was never completed. *Star Trek II: The Wrath of Khan* premiered a short CGI sequence called The Genesis Wave in June 1982. The first two films to make heavy investments in Solid 3D CGI, *Tron* (1982) and *The Last Starfighter* (1984), were commercial failures, causing most directors to relegate CGI to images that were supposed to look like they were created by a computer.

Another Star Trek movie - 1986's *Star Trek IV: The Voyage Home* - features a little-known use of CGI that would prove to be a major stepping stone for the technology in years to come. The dream-like sequence where Kirk and his crew (in a commandeered Klingon starship) travel back in time features computer-generated images of the crew's faces "morphing" into one another. Although the images look like clay sculptures, they represent the antecedents of the photo-realistic morphing that would become popular in the early 90s, and the ILM team responsible would be sought out by James Cameron three years later (see below).

The first movie to feature photo-realistic CGI integrated seamlessly into scenes was *The Abyss* (1989). A five minute sequence featuring an animated water tentacle or "pseudopod" was created using groundbreaking new algorithms design by ILM, and supervised by Denis Muren for James Cameron's underwater action movie. It featured reflection, refraction and a short "morphing" sequence, taking on the facial forms of actors Ed Harris and Mary Elizabeth Mastrantonio. Although Cameron's subsequent movie, *Terminator 2: Judgment Day*, would be the one that brought CGI to the attention of the masses, it was The Abyss that represented the true technological milestone, and the foundation upon which the modern CGI industry would be built.

1991 could be considered the breakout year for the new CGI technology. Two huge hits, *Terminator 2: Judgment Day* and *Beauty and the Beast*, both made heavy use of CGI.

These successes marked Hollywood's transition from stop-motion animation and conventional optical effects to digital techniques. *Beauty and the Beast* became the first animated film to be nominated for Best Picture, and *Terminator 2: Judgment Day* won the Oscar for Best Visual Effects.

In 1993, another affirmation of CGI came from *Jurassic Park* where CGI dinosaurs were integrated into hydraulically-controlled life-sized puppets shot on-set. The palette of tools available from CGI was becoming larger. In 1994, CGI was used to create the special effects for *Forrest Gump*. The most noteworthy effects shots were those that featured the digital removal of actor Gary Sinise's legs. Other effects included a napalm strike, the fast-moving Ping-Pong balls, and the digital insertion of Tom Hanks into several scenes of historical footage.

In 1993, *Babylon 5* and *SeaQuest* became the first television series to use CGI as the primary method for their visual effects (rather than using hand-built models). It also marked the first TV use of virtual sets. That same year, *Insektors* became the first full-length completely computer animated TV series . Soon after, in 1994, the hit Canadian CGI show *ReBoot* aired.

In 1995, the first fully computer-generated feature film, Disney-Pixar's *Toy Story*, was a resounding commercial success. Additional digital animation studios such as Blue Sky Studios (20th Century Fox), DNA Productions (Paramount Pictures and Warner Bros.), Omation Studios (Paramount Pictures), Sony Pictures Animation (Columbia Pictures), Vanguard Animation (Walt Disney Pictures, Lions Gate Entertainment and 20th Century Fox), Big Idea Productions (Universal Pictures and FHE Pictures), Animal Logic (Warner Bros.) and Pacific Data Images (Dreamworks SKG) went into production, and existing animation companies, such as The Walt Disney Company, began to make a transition from traditional animation to CGI. Between 1995 and 2005 the average effects budget for a wide-release feature film skyrocketed from $5 million to $40 million. According to one studio executive, as of 2005, more than half of feature films have significant effects. However, CGI has made up for the expenditures by grossing over 20% more than their real-life counterparts.

In the early 2000s, computer-generated imagery became the dominant form of special effects. The technology progressed to the point that it became possible to include virtual stunt doubles. Camera tracking software was refined to allow increasingly complex visual effects developments that were previously impossible. Computer-generated extras also became used extensively in crowd scenes with advanced flocking and crowd simulation software. Virtual sets, in which part or all of the background of a shot is digitally generated, also became commonplace. The timeline of CGI in film and television shows a detailed list of pioneering uses of computer-generated imagery in film and television.

CGI for films is usually rendered at about 1.4–6 megapixels. *Toy Story*, for example, was rendered at 1536 × 922 (1.42MP). The time to render one frame is typically around 2–3 hours, with ten times that for the most complex scenes. This time has not changed much in the last decade, as image quality has progressed at the same rate as improvements in hardware, since with faster machines, more and more complexity becomes feasible. Exponential increases in GPUs processing power, as well as massive increases in parallel CPU power, storage and memory speed and size have greatly increased CGI's potential.

In 2001, Square Pictures created the CGI film *Final Fantasy: The Spirits Within*, which made headlines for attempting to create photo-realistic human actors. The film was not a box-office success. Some commentators have suggested this may be partly because the lead CGI characters had facial features which fell into the uncanny valley. Square Pictures produced only two more films using a similar visual style *Final Flight of the Osiris*, a short film which served as a prologue to *The Matrix Reloaded* and *Final Fantasy VII: Advent Children*, based on their extremely popular video game series.

Developments in CGI technologies are reported each year at SIGGRAPH, an annual conference on computer graphics and interactive techniques, attended each year by tens of thousands of computer professionals. Developers of computer games and 3D video cards strive to achieve the same visual quality on personal computers in real-time as is possible for CGI films and animation. With the rapid advancement of real-time rendering quality, artists began to use game engines to render non-interactive movies. This art form is called *machinima*.

Methods of animating virtual characters

In most 3D computer animation systems, an animator creates a simplified representation of a character's anatomy, analogous to a skeleton or stick figure. The position of each segment of the skeletal model is defined by *animation variables*, or Avars. In human and animal characters, many parts of the skeletal model correspond to actual bones, but skeletal animation is also used to animate other things, such as facial features (though other methods for facial animation exist). The character "Woody" in *Toy Story*, for example, uses 700 Avars, including 100 Avars in the face. The computer does not usually render the skeletal model directly (it is invisible), but uses the skeletal model to compute the exact position and orientation of the character, which is eventually rendered into an image. Thus by changing the values of Avars over time, the animator creates motion by making the character move from frame to frame.

There are several methods for generating the Avar values to obtain realistic motion. Traditionally, animators manipulate the Avars directly. Rather than set Avars for every frame, they usually set Avars at strategic points (frames) in time and let the computer interpolate or 'tween' between them, a process called keyframing. Keyframing puts control in the hands of the animator, and has roots in hand-drawn traditional animation.

In contrast, a newer method called motion capture makes use of live action. When computer animation is driven by motion capture, a real performer acts out the scene as if they were the character to be animated. His or her motion is recorded to a computer using video cameras and markers, and that performance is then applied to the animated character.

Each method has its advantages, and as of 2007, games and films are using either or both of these methods in productions. Keyframe animation can produce motions that would be difficult or impossible to act out, while motion capture can reproduce the subtleties of a particular actor. For example, in the 2006 film Pirates of the Caribbean: Dead Man's Chest, actor Bill Nighy provided the performance for the character Davy Jones. Even though Nighy himself doesn't appear in the film, the movie benefited from his performance by recording the nuances of his body language, posture, facial expressions, etc. Thus motion capture is appropriate in situations where believable, realistic behavior and action is required, but the types of characters required exceed what can be done through conventional costuming.

Creating characters and objects on a computer

3D **computer animation** combines 3D models of objects and programmed or hand "keyframed" movement. Models are constructed out of geometrical vertices, faces, and edges in a 3D coordinate system. Objects are sculpted much like real clay or plaster, working from general forms to specific details with various sculpting tools. A bone/joint animation system is set up to deform the CGI model (e.g., to make a humanoid model walk). In a process called rigging, the virtual marionette is given various controllers and handles for controlling movement. Animation data can be created using motion capture, or keyframing by a human animator, or a combination of the two.

3D models rigged for animation may contain thousands of control points - for example, the character "Woody" in Pixar's movie *Toy Story*, uses 700 specialized animation controllers. Rhythm and Hues Studios labored for two years to create Aslan in the movie *The Chronicles of Narnia: The Lion, the Witch and the Wardrobe* which had about 1851 controllers, 742 in just the face alone. In the 2004 film *The Day After Tomorrow*, designers had to design forces of extreme weather with the help of video references and accurate meteorological facts. For the 2005 remake of *King Kong*, actor Andy Serkis was used to help designers pinpoint the gorilla's prime location in the shots and used his expressions to model "human" characteristics onto the creature. Serkis had earlier provided the voice and performance for Gollum in J. R. R. Tolkien's *The Lord of the Rings* trilogy.

Computer animation development equipment

Computer animation can be created with a computer and animation software. Some impressive animation can be achieved even with basic programs; however the rendering can take a lot of time on an ordinary home computer. Because of this, video game animators tend to use low resolution, low polygon count renders, such that the graphics can be rendered in real time on a home computer. Photorealistic animation would be impractical in this context.

Professional animators of movies, television, and video sequences on computer games make photorealistic animation with high detail. This level of quality for movie animation would take tens to hundreds of years to create on a home computer. Many powerful workstation computers are used instead. Graphics workstation computers use two to four processors, and thus are a lot more powerful than a home computer, and are specialized for rendering. A large number of workstations (known as a render farm) are networked together to effectively act as a giant computer. The result is a computer-animated movie that can be completed in about one to five years (this process is not comprised solely of rendering, however). A workstation typically costs $2,000 to $16,000, with the more expensive stations being able to render much faster, due to the more technologically advanced hardware that they contain. Pixar's Renderman is rendering software which is widely used as the movie animation industry standard, in competition with Mental Ray. It can be bought at the official Pixar website for about $3,500. It will work on Linux, Mac OS X, and Microsoft Windows based graphics workstations along with an animation program such as Maya and Softimage XSI. Professionals also use digital movie cameras, motion capture or performance capture, bluescreens, film editing software, props, and other tools for movie animation.

The future

One open challenge in computer animation is a photorealistic animation of humans. Currently, most computer-animated movies show animal characters (*A Bug's Life*, *Finding Nemo*, *Ratatouille*, *Ice Age*, *Over the Hedge*), fantasy characters (*Monsters Inc.*, *Shrek*, *Teenage Mutant Ninja Turtles 4*, *Monsters vs. Aliens*), anthropomorphic machines (*Cars*, *WALL-E*, *Robots*) or cartoon-like humans (*The Incredibles*, *Despicable Me*, *Up*). The movie *Final Fantasy: The Spirits Within* is often cited as the first computer-generated movie to attempt to show realistic-looking humans. However, due to the enormous complexity of the human body, human motion, and human biomechanics, realistic simulation of humans remains largely an open problem. Another problem is the distasteful psychological response to viewing nearly perfect animation of humans, known as "the uncanny valley." It is one of the "holy grails" of computer animation. Eventually, the goal is to create software where the animator can generate a movie sequence showing a photorealistic human character, undergoing physically-plausible motion, together with clothes, photorealistic hair, a complicated natural background, and possibly interacting with other simulated human characters. This could be done in a way that the viewer is no longer able to tell if a particular movie sequence is computer-generated, or created using real actors in front of movie cameras. Complete human realism is not likely to happen very soon, but when it does it may have major repercussions for the film industry.

For the moment it looks like three dimensional computer animation can be divided into two main directions; photorealistic and non-photorealistic rendering. Photorealistic computer animation can itself be divided into two subcategories; real photorealism (where performance capture is used in the creation of the virtual human characters) and stylized photorealism. Real photorealism is what Final Fantasy tried to achieve and will in the future most likely have the ability to give us live action fantasy features as *The Dark Crystal* without having to use advanced puppetry and animatronics, while *Antz* is an example on stylistic photorealism (in the future stylized photorealism will be able to replace traditional stop motion animation as in *Corpse Bride*). None of these mentioned are perfected as of yet, but the progress continues.

The non-photorealistic/cartoonish direction is more like an extension of traditional animation, an attempt to make the animation look like a three dimensional version of a cartoon, still using and perfecting the main principles of animation articulated by the Nine Old Men, such as squash and stretch.

While a single frame from a photorealistic computer-animated feature will look like a photo if done right, a single frame vector from a cartoonish computer-animated feature will look like a painting (not to be confused with cel shading, which produces an even simpler look).

Detailed examples and pseudocode

In 2D computer animation, moving objects are often referred to as “sprites.” A sprite is an image that has a location associated with it. The location of the sprite is changed slightly, between each displayed frame, to make the sprite appear to move. The following pseudocode makes a sprite move from left to right:

```
var int x := 0, y := screenHeight / 2;
```

```
while x < screenWidth
drawBackground()
drawSpriteAtXY (x, y) // draw on top of the background
x := x + 5 // move to the right
```

Computer animation uses different techniques to produce animations. Most frequently, sophisticated mathematics is used to manipulate complex three dimensional polygons, apply "textures", lighting and other effects to the polygons and finally rendering the complete image. A sophisticated graphical user interface may be used to create the animation and arrange its choreography. Another technique called constructive solid geometry defines objects by conducting boolean operations on regular shapes, and has the advantage that animations may be accurately produced at any resolution.

Let's step through the rendering of a simple image of a room with flat wood walls with a grey pyramid in the center of the room. The pyramid will have a spotlight shining on it. Each wall, the floor and the ceiling is a simple polygon, in this case, a rectangle. Each corner of the rectangles is defined by three values referred to as X, Y and Z. X is how far left and right the point is. Y is how far up and down the point is, and Z is far in and out of the screen the point is. The wall nearest us would be defined by four points: (in the order x, y, z). Below is a representation of how the wall is defined

```
(0, 10, 0)                        (10, 10, 0)

(0,0,0)                           (10, 0, 0)
```

The far wall would be:

```
(0, 10, 20)                        (10, 10, 20)

(0, 0, 20)                         (10, 0, 20)
```

The pyramid is made up of five polygons: the rectangular base, and four triangular sides. To draw this image the computer uses math to calculate how to project this image, defined by three dimensional data, onto a two dimensional computer screen.

First we must also define where our view point is, that is, from what vantage point will the scene be drawn. Our view point is inside the room a bit above the floor, directly in front of the pyramid. First the computer will calculate which polygons are visible. The near wall will not be displayed at all, as it is behind our view point. The far side of the pyramid will also not be drawn as it is hidden by the front of the pyramid.

Next each point is perspective projected onto the screen. The portions of the walls 'furthest' from the view point will appear to be shorter than the nearer areas due to perspective. To make the walls look like wood, a wood pattern, called a texture, will be drawn on them. To accomplish this, a technique called "texture mapping" is often used. A small drawing of wood that can be repeatedly drawn in a matching tiled pattern (like wallpaper) is stretched and drawn onto the walls' final shape. The pyramid is solid grey so its surfaces can just be rendered as grey. But we also have a spotlight. Where its light falls we lighten colors, where objects blocks the light we darken colors.

Next we render the complete scene on the computer screen. If the numbers describing the position of the pyramid were changed and this process repeated, the pyramid would appear to move.

Movies

CGI short films have been produced as independent animation since 1976, though the popularity of computer animation (especially in the field of special effects) skyrocketed during the modern era of U.S. animation. The first completely computer-generated television series was *ReBoot*, in 1994, and the first completely computer-generated animated movie was *Toy Story* (1995).

Amateur animation

The popularity of websites which allows members to upload their own movies for others to view has created a growing community of amateur computer animators. With utilities and programs often included free with modern operating systems, many users can make their own animated movies and shorts. Several free and open source animation software applications exist as well. A popular amateur approach to animation is via the animated GIF format, which can be uploaded and seen on the web easily.

Chapter- 5

Other Applications of Motion Planning

Robotic mapping (Application of Motion Planning)

Robotic mapping can be used for serving robot guide

The problem of **Robotic mapping** is related to cartography. The goal is for an autonomous robot to be able to construct (or use) a map or floor plan and to localize itself in it.

Robotic mapping is that branch of one, which deals with the study and application of ability to construct map or floor plan by the autonomous robot and to localize itself in it.

Todd et al. (1994) have shown that evolutionarily shaped blind action may suffice to keep some animals alive. For some insects for example, the environment is not interpreted as a map, and they survive only with a triggered response.

But a slightly more elaborated navigation strategy dramatically enhances the capabilities of the robot. Cognitive maps (Tolman 1948) enable planning capacities, and use of current perceptions, memorized events, and expected consequences.

The problem can be decomposed in three processes (Levitt and Lawton 1990 ; Balakrishnan et al. 1999) : map learning, localization, path-planning.

Available information

The robot has two sources of information: the idiothetic and the allothetic sources.

When in motion, a robot can use dead reckoning methods such as tracking the number of revolutions of its wheels; this corresponds to the idiothetic source and can give the absolute position of the robot, but it is subject to cumulative error which can grow quickly.

The allothetic source corresponds to the sensors of the robot, like a camera, a microphone, laser or sonar. The problem here is "perceptual aliasing". This means that two different places can be perceived as the same. For example, in a building, it may be impossible for to determine your location solely with the visual information, because all the corridors may look the same.

Map representation

The internal representation of the map can be "metric" or "topological":

- The metric framework is the most common for humans and considers a two dimensional space in which it places the objects. The objects are placed with precise coordinates. This representation is very useful, but is sensitive to noise and it is difficult to calculate precisely the distances.

- The topological framework only considers places and relations between them. Often, the distances between places are stored. The map is then a graph, in which the nodes corresponds to places and arcs correspond to the paths.

Many techniques use probabilistic representations of the map, in order to handle uncertainty.

There are three main methods of Map representations:

Free Space Maps

- Spatial graphs
- Generalized Voronoi Diagrams

Quadtree

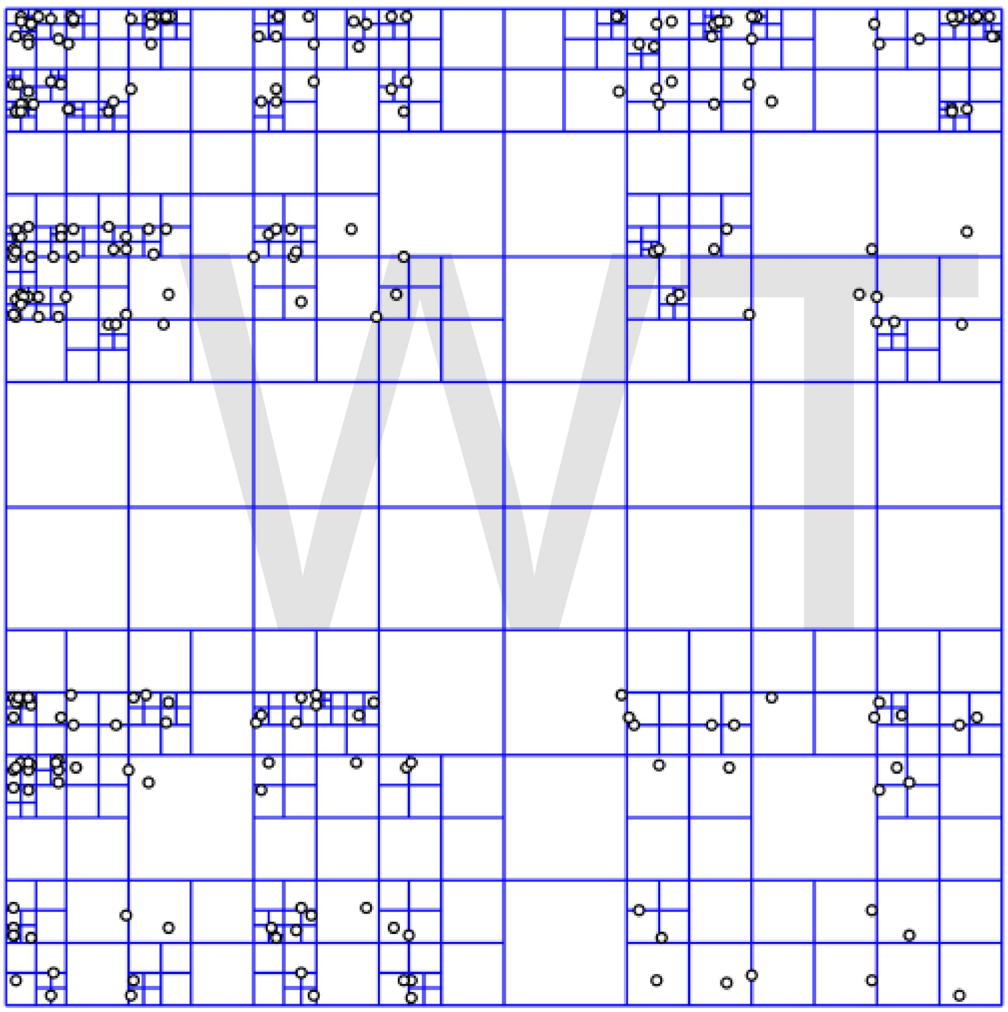

A region quadtree with point data

A **quadtree** is a tree data structure in which each internal node has exactly four children. Quadtrees are most often used to partition a two dimensional space by recursively subdividing it into four quadrants or regions. The regions may be square or rectangular, or may have arbitrary shapes. This data structure was named a quadtree by Raphael Finkel and J.L. Bentley in 1974. A similar partitioning is also known as a *Q-tree*. All forms of Quadtrees share some common features:

- They decompose space into adaptable cells
- Each cell (or bucket) has a maximum capacity. When maximum capacity is reached, the bucket splits
- The tree directory follows the spatial decomposition of the Quadtree.

Types

Quadtrees may be classified according to the type of data they represent, including areas, points, lines and curves. Quadtrees may also be classified by whether the shape of the tree is independent of the order data is processed. Some common types of quadtrees are:

The region quadtree

The region quadtree represents a partition of space in two dimensions by decomposing the region into four equal quadrants, subquadrants, and so on with each leaf node containing data corresponding to a specific subregion. Each node in the tree either has exactly four children, or has no children (a leaf node). The region quadtree is not strictly a 'tree' - as the positions of subdivisions are independent of the data. They are more precisely called 'tries'.

A region quadtree with a depth of n may be used to represent an image consisting of $2^n \times 2^n$ pixels, where each pixel value is 0 or 1. The root node represents the entire image region. If the pixels in any region are not entirely 0s or 1s, it is subdivided. In this application, each leaf node represents a block of pixels that are all 0s or all 1s.

A region quadtree may also be used as a variable resolution representation of a data field. For example, the temperatures in an area may be stored as a quadtree, with each leaf node storing the average temperature over the subregion it represents.

If a region quadtree is used to represent a set of point data (such as the latitude and longitude of a set of cities), regions are subdivided until each leaf contains at most a single point.

Point quadtree

The point quadtree is an adaptation of a binary tree used to represent two dimensional point data. It shares the features of all quadtrees but is a true tree as the center of a subdivision is always on a point. The tree shape depends on the order data is processed. It is often very efficient in comparing two dimensional ordered data points, usually operating in O(log n) time.

Node structure for a point quadtree

A node of a point quadtree is similar to a node of a binary tree, with the major difference being that it has four pointers (one for each quadrant) instead of two ("left" and "right") as in an ordinary binary tree. Also a key is usually decomposed into two parts, referring to x and y coordinates. Therefore a node contains following information:

- 4 Pointers: quad['NW'], quad['NE'], quad['SW'], and quad['SE']
- point; which in turn contains:
    - key; usually expressed as x, y coordinates
    - value; for example a name

Edge quadtree

Edge quadtree are specifically used to store lines rather than points. Curves are approximated by subdividing cells to a very fine resolution. This can result in extremely unbalanced trees which may defeat the purpose of indexing.

Some common uses of quadtrees

- Image representation

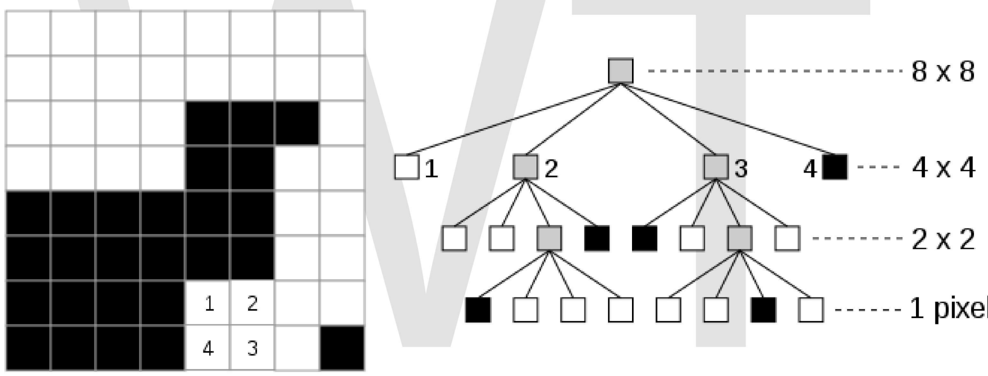

- Spatial indexing
- Efficient collision detection in two dimensions
- View frustum culling of terrain data
- Storing sparse data, such as a formatting information for a spreadsheet or for some matrix calculations
- Solution of multidimensional fields (computational fluid dynamics, electromagnetism)
- Conway's Game of Life simulation program.
- State estimation

Quadtrees are the two-dimensional analog of octrees.

These employ the notion of a grid, but permit the resolution of the grid to vary so that it can become finer where more accuracy is needed and more coarse where the map is uniform.

Map learning

Map-learning cannot be separated from the localization process, and a difficulty arises when errors in localization are incorporated into the map. This problem is commonly referred to as Simultaneous localization and mapping (SLAM).

An important additional problem is to determine whether the robot is in a part of environment already stored or never visited. One way to solve this problem is by using electric beacons.

Path planning

Path planning is an important issue as it allows a robot to get from point A to point B. Path planning algorithms are measured by their computational complexity. The feasibility of real-time motion planning is dependent on the accuracy of the map (or floorplan), on robot localization and on the number of obstacles. Topologically, the problem of path planning is related to the shortest path problem problem of finding a route between two nodes in a graph.

Robot navigation

Outdoor robots can use GPS in a similar way to automotive navigation systems.

Alternative systems can be used with floor plan instead of maps for indoor robots, combined with localization wireless hardware.

Electric beacons also have been proposed for cheap robot navigational systems.

Computer-aided architectural design

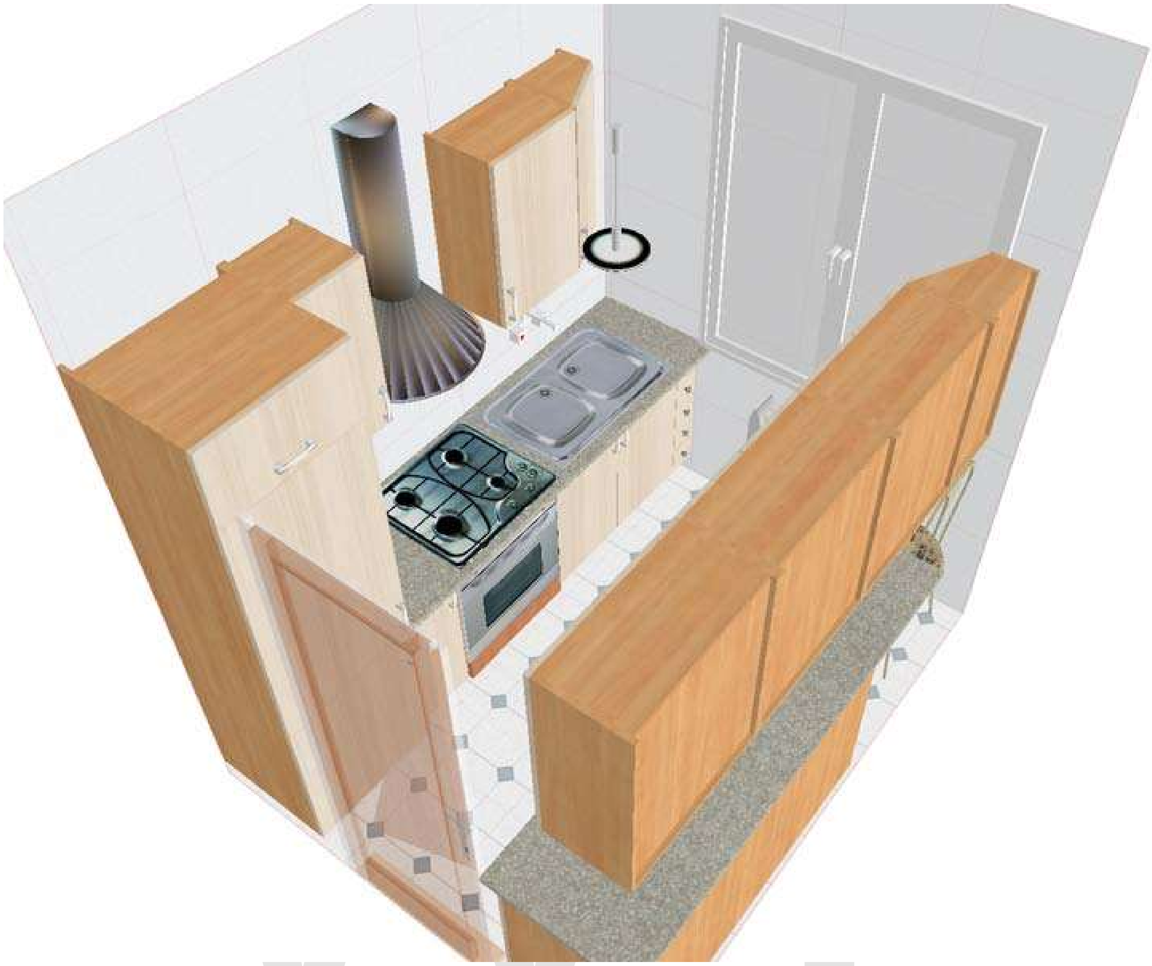

Example of Computer-aided architectural design

Computer-aided architectural design (CAAD) software programs are the repository of accurate and comprehensive records of buildings and are used by architects and architectural companies.

The first program was installed back in the 1960's, to help architects save time instead of drawing their blueprints. Computer-aided design also known as CAD was originally the type of program that architects used, but since CAD couldn't offer all the tools that architects needed to complete a project, CAAD developed as a distinct class of software.

Overview

All CAD and CAAD systems employ a database with geometric and other properties of objects; they all have some kind of graphic user interface to manipulate a visual representation rather than the database; and they are all more or less concerned with assembling designs from standard and non-standard pieces. Currently, the main distinction which causes one to speak of CAAD rather than CAD lies in the domain knowledge (architecture-specific objects, techniques, data, and process support) embedded in the system. A CAAD system differs from other CAD systems in two respects:

- It has an explicit object database of building parts and construction knowledge.
- It explicitly supports the creation of architectural objects.

In a more general sense, CAAD also refers to the use of any computational technique in the field of architectural design other than by means of architecture-specific software. For example, software which is specifically developed for the computer animation industry (e.g. Maya and 3DStudio Max), is also used in architectural design. The exact distinction of what properly belongs to CAAD is not always clear. Specialized software, for example for calculating structures by means of the finite element method, is used in architectural design and in that sense may fall under CAAD. On the other hand, such software is seldom used to create new designs.

In 1974 Caad became a current word and was a common topic of commercial modernization.

Three dimensional objects

CAAD has two types of structures in its program. The first system is surface structure which provides a graphics medium to represent three dimensional objects using two dimensional representations. Also algorithms that allow the generation of patterns and their analysis using programmed criteria, and data banks that store information about the problem at hand and the standards and regulations that applies to it. The second system is deep structure which means that the operations performed by the computer have natural limitations. Computer hardware and machine languages that are supported by these make it easy to perform arithmetical operations quickly and accurately. Also almost an illogical number of layers of symbolic processing can be built enabling the functionalities that are found at the surface.

Advantages

Another advantage to CAAD is the two way mapping of activities and functionalities.NOS. The two instances of mapping are indicated to be between the surface structures (TM1) and the deep structures (TM2). These mappings are abstractions that are introduced in order to discuss the process of design and deployment of CAAD systems. In designing the systems the system developers usually consider TM1. Here a one-to-one mapping is the typical statement, which is to develop a computer based functionality that maps as closely as possible into a corresponding manual design activity, for example, drafting of stairs, checking spatial conflict between building systems, and generating perspectives from orthogonal views. The architectural design processes tend to integrate models isolated so far. Many different kinds of expert knowledge, tools, visualization techniques, and media are to be combined. The design process covers the complete life cycle of the building. The areas that are covered are construction, operations, reorganization, as well as destruction. Considering the shared use of digital design tools and the exchange of information and knowledge between designers and across different projects, we speak of a design continuum.

An architect's work involves mostly visually represented data. Problems are often outlined and dealt with in a graphical approach. Only this form of expression serves as a basis for work and discussion. Therefore, the designer should have a maximum visual control over the processes taking place with in the design continuum. Further questions occur about navigation, associative information access, programming and communication with in very large data sets.

Chapter- 6

Systems of Motion Planning in Robotics

Robotics

The Shadow robot hand system

A Pick and Place robot in a factory

Robotics is the engineering science and technology of robots, and their design, manufacture, application, and structural disposition. Robotics is related to electronics, mechanics, and software. The word "robot" was introduced to the public by Czech writer Karel Čapek in his play R.U.R. (Rossum's Universal Robots), published in 1920. The term "robotics" was coined by Isaac Asimov in his 1941 science fiction short-story "Liar!"

History

Stories of artificial helpers and companions and attempts to create them have a long history.

In 1921, Czech writer Karel Čapek introduced the word "robot" in his play *R.U.R. (Rossum's Universal Robots)*. The word "robot" comes from the word "robota", meaning, in Czech, "forced labour, drudgery".

In 1927, the Maschinenmensch ("machine-human"), a gynoid humanoid robot, also called "Parody", "Futura", "Robotrix", or the "Maria impersonator" (played by German actress Brigitte Helm), the first and perhaps the most memorable depiction of a robot ever to appear on film, was depicted in Fritz Lang's film Metropolis.

In 1942, Isaac Asimov formulatead the Three Laws of Robotics, and in the process of doing so, coined the word "robotics".

In 1948, Norbert Wiener formulated the principles of cybernetics, the basis of practical robotics.

Fully autonomous robots only appeared in the second half of the 20th century. The first digitally operated and programmable robot, the Unimate, was installed in 1961 to lift hot pieces of metal from a die casting machine and stack them. Today, commercial and industrial robots are in widespread use performing jobs more cheaply or more accurately and reliably than humans. They are also employed in jobs which are too dirty, dangerous, or dull to be suitable for humans. Robots are widely used in manufacturing, assembly, and packing; transport; earth and space exploration; surgery; weaponry; laboratory research; safety; and mass production of consumer and industrial goods.

Etymology

According to the *Oxford English Dictionary*, the word *robotics* was first used in print by Isaac Asimov, in his science fiction short story "Liar!", published in May 1941 in *Astounding Science Fiction*. Asimov was unaware that he was coining the term; since the science and technology of electrical devices is *electronics*, he assumed *robotics* already referred to the science and technology of robots. However, in some of Asimov's other works, he states that the first use of the word *robotics* was in his short story *Runaround* (Astounding Science Fiction, March 1942). The word *robotics* was derived from the word *robot*, which was introduced to the public by Czech writer Karel Čapek in his play *R.U.R. (Rossum's Universal Robots)*, which premiered in 1921.

Components

Structure

The structure of a robot is usually mostly mechanical and can be called a kinematic chain (its functionality being similar to the skeleton of the human body). The chain is formed of links (its bones), actuators (its muscles), and joints which can allow one or more degrees of freedom. Most contemporary robots use open serial chains in which each link connects the one before to the one after it. These robots are called serial robots and often resemble the human arm. Some robots, such as the Stewart platform, use a closed parallel kinematical chain. Other structures, such as those that mimic the mechanical structure of humans, various animals, and insects, are comparatively rare. However, the development and use of such structures in robots is an active area of research (e.g. biomechanics). Robots used as manipulators have an end effector mounted on the last link. This end effector can be anything from a welding device to a mechanical hand used to manipulate the environment.

Power source

At present; mostly (lead-acid) batteries are used, but potential power sources could be:

- pneumatic (compressed gases)
- hydraulics (compressed liquids)
- flywheel energy storage
- organic garbage (through anaerobic digestion)
- faeces (human, animal); may be interesting in a military context as feces of small combat groups may be reused for the energy requirements of the robot assistant
- still untested energy sources (e.g. Nuclear Fusion reactors, ...)
- radioactive source (such as with the proposed Ford car of the '50s); to those proposed in movies such as *Red Planet*

Actuation

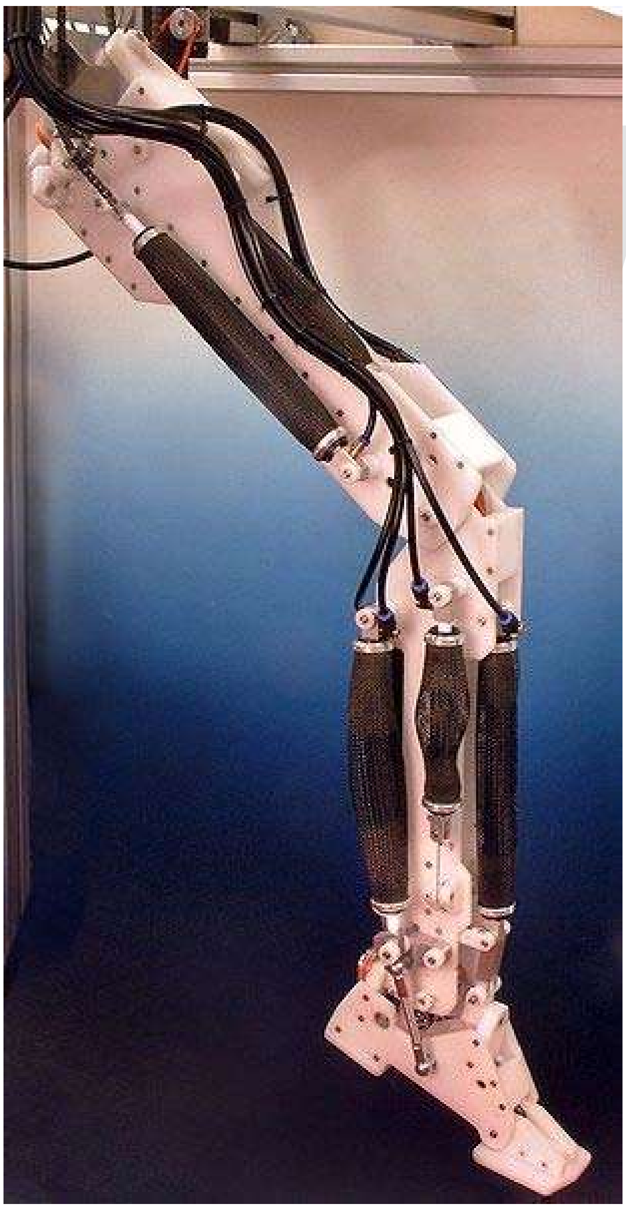

A robot leg powered by Air Muscles

Actuators are like the "muscles" of a robot, the parts which convert stored energy into movement. By far the most popular actuators are electric motors that spin a wheel or gear, and linear actuators that control industrial robots in factories. But there are some recent advances in alternative types of actuators, powered by electricity, chemicals, or compressed air:

- **Electric motors:** The vast majority of robots use electric motors, often brushed and brushless DC motors in portable robots or AC motors in industrial robots and CNC machines.
- **Linear Actuators:** Various types of linear actuators move in and out instead of by spinning, particularly when very large forces are needed such as with industrial robotics. They are typically powered by compressed air (pneumatic actuator) or an oil (hydraulic actuator).
- **Series Elastic Actuators:** A spring can be designed as part of the motor actuator, to allow improved force control. It has been used in various robots, particularly walking humanoid robots.
- **Air muscles:** (Also known as Pneumatic Artificial Muscles) are special tubes that contract (typically up to 40%) when air is forced inside it. They have been used for some robot applications.
- **Muscle wire:** (Also known as Shape Memory Alloy, Nitinol or Flexinol Wire) is a material that contracts slightly (typically under 5%) when electricity runs through it. They have been used for some small robot applications.
- **Electroactive Polymers:** (EAPs or EPAMs) are a new plastic material that can contract substantially (up to 400%) from electricity, and have been used in facial muscles and arms of humanoid robots, and to allow new robots to float, fly, swim or walk.
- **Piezo motor:** A recent alternative to DC motors are piezo motors or ultrasonic motors. These work on a fundamentally different principle, whereby tiny piezoceramic elements, vibrating many thousands of times per second, cause linear or rotary motion. There are different mechanisms of operation; one type uses the vibration of the piezo elements to walk the motor in a circle or a straight line. Another type uses the piezo elements to cause a nut to vibrate and drive a screw. The advantages of these motors are nanometer resolution, speed, and available force for their size. These motors are already available commercially, and being used on some robots.
- **Elastic nanotubes:** These are a promising, early-stage experimental technology. The absence of defects in nanotubes enables these filaments to deform elastically by several percent, with energy storage levels of perhaps 10 J/cm^3 for metal nanotubes. Human biceps could be replaced with an 8 mm diameter wire of this material. Such compact "muscle" might allow future robots to outrun and outjump humans.

Sensing

Touch

Current robotic and prosthetic hands receive far less tactile information than the human hand. Recent research has developed a tactile sensor array that mimics the mechanical properties and touch receptors of human fingertips. The sensor array is constructed as a rigid core surrounded by conductive fluid contained by an elastomeric skin. Electrodes are mounted on the surface of the rigid core and are connected to an impedance-measuring device within the core. When the artificial skin touches an object the fluid path around the electrodes is deformed, producing impedance changes that map the forces received from the object. The researchers expect that an important function of such artificial fingertips will be adjusting robotic grip on held objects.

In 2009, scientists from several European countries and Israel developed a prosthetic hand, called SmartHand, which functions like a real one, allowing patients to write with it, type on a keyboard, play piano and perform other fine movements. The prosthesis has sensors which enable the patient to sense real feeling in its fingertips.

Vision

Computer vision is the science and technology of machines that see. As a scientific discipline, computer vision is concerned with the theory behind artificial systems that extract information from images. The image data can take many forms, such as video sequences and views from cameras.

In most practical computer vision applications, the computers are pre-programmed to solve a particular task, but methods based on learning are now becoming increasingly common.

Computer vision systems rely on image sensors which detect electromagnetic radiation which is typically in the form of either visible light or infra-red light. The sensors are designed using solid-state physics. The process by which light propagates and reflects off surfaces is explained using optics. Sophisticated image sensors even require quantum mechanics to provide a complete understanding of the image formation process.

There is a subfield within computer vision where artificial systems are designed to mimic the processing and behavior of biological systems, at different levels of complexity. Also, some of the learning-based methods developed within computer vision have their background in biology.

Manipulation

Robots which must work in the real world require some way to manipulate objects; pick up, modify, destroy, or otherwise have an effect. Thus the 'hands' of a robot are often referred to as end effectors, while the arm is referred to as a manipulator. Most robot arms have replaceable effectors, each allowing them to perform some small range of tasks. Some have a fixed manipulator which cannot be replaced, while a few have one very general purpose manipulator, for example a humanoid hand.

- **Mechanical Grippers:** One of the most common effectors is the gripper. In its simplest manifestation it consists of just two fingers which can open and close to pick up and let go of a range of small objects. Fingers can for example be made of a chain with a metal wire run trough it.
- **Vacuum Grippers:** Pick and place robots for electronic components and for large objects like car windscreens, will often use very simple vacuum grippers. These are very simple astrictive devices, but can hold very large loads provided the prehension surface is smooth enough to ensure suction.
- **General purpose effectors:** Some advanced robots are beginning to use fully humanoid hands, like the Shadow Hand, MANUS, and the Schunk hand. These highly dexterous manipulators, with as many as 20 degrees of freedom and hundreds of tactile sensors.

For the definitive guide to all forms of robot endeffectors, their design, and usage consult the book "Robot Grippers".

Locomotion

Rolling robots

Segway in the Robot museum in Nagoya.

For simplicity, most mobile robots have four wheels. However, some researchers have tried to create more complex wheeled robots, with only one or two wheels. These can have certain advantages such as greater efficiency, reduced parts, and allow a robot to navigate in tight places that a four wheeled robot would not be able to.

- **Two-wheeled balancing:** Balancing robots generally use a gyroscope to detect how much a robot is falling and then drive the wheels proportionally in the opposite direction, to counter-balance the fall at hundreds of times per second, based on the dynamics of an inverted pendulum. Many different balancing robots have been designed. While the Segway is not commonly thought of as a robot, it can be thought of as a component of a robot, such as NASA's Robonaut that has been mounted on a Segway.

- **One-wheeled balancing:** A one-wheeled balancing robot is an extension of a two-wheeled balancing robot so that it can move in any 2D direction using a round ball as its only wheel. Several one-wheeled balancing robots have been designed recently, such as Carnegie Mellon University's "Ballbot" that is the approximate height and width of a person, and Tohoku Gakuin University's "BallIP". Because of the long, thin shape and ability to maneuver in tight spaces, they have the potential to function better than other robots in environments with people.
- **Spherical orb robots:** Several attempts have been made in robots that are completely inside a spherical ball, either by spinning a weight inside the ball, or by rotating the outer shells of the sphere. These have also been referred to as an orb bot or a ball bot
- **Six-wheeled robots:** Using six wheels instead of four wheels can give better traction or grip in outdoor terrain such as on rocky dirt or grass.
- **Tracked robots:** Tank tracks provide even more traction than a six-wheeled robot. Tracked wheels behave as if they were made of hundreds of wheels, therefore are very common for outdoor and military robots, where the robot must drive on very rough terrain. However, they are difficult to use indoors such as on carpets and smooth floors. Examples include NASA's Urban Robot "Urbie".

Walking robots

iCub robot, designed by the RobotCub Consortium

Walking is a difficult and dynamic problem to solve. Several robots have been made which can walk reliably on two legs, however none have yet been made which are as robust as a human. Many other robots have been built that walk on more than two legs, due to these robots being significantly easier to construct. Hybrids too have been proposed in movies such as I, Robot, where they walk on 2 legs and switch to 4 (arms+legs) when going to a sprint. Typically, robots on 2 legs can walk well on flat floors, and can occasionally walk up stairs. None can walk over rocky, uneven terrain. Some of the methods which have been tried are:

- **ZMP Technique:** The Zero Moment Point (ZMP) is the algorithm used by robots such as Honda's ASIMO. The robot's onboard computer tries to keep the total inertial forces (the combination of earth's gravity and the acceleration and deceleration of walking), exactly opposed by the floor reaction force (the force of the floor pushing back on the robot's foot). In this way, the two forces cancel out, leaving no moment (force causing the robot to rotate and fall over). However, this is not exactly how a human walks, and the difference is obvious to human observers, some of whom have pointed out that ASIMO walks as if it needs the lavatory. ASIMO's walking algorithm is not static, and some dynamic balancing is used (See below). However, it still requires a smooth surface to walk on.
- **Hopping:** Several robots, built in the 1980s by Marc Raibert at the MIT Leg Laboratory, successfully demonstrated very dynamic walking. Initially, a robot with only one leg, and a very small foot, could stay upright simply by hopping. The movement is the same as that of a person on a pogo stick. As the robot falls to one side, it would jump slightly in that direction, in order to catch itself. Soon, the algorithm was generalised to two and four legs. A bipedal robot was demonstrated running and even performing somersaults. A quadruped was also demonstrated which could trot, run, pace, and bound.
- **Dynamic Balancing** or controlled falling: A more advanced way for a robot to walk is by using a dynamic balancing algorithm, which is potentially more robust than the Zero Moment Point technique, as it constantly monitors the robot's motion, and places the feet in order to maintain stability. This technique was recently demonstrated by Anybots' Dexter Robot, which is so stable, it can even jump. Another example is the TU Delft Flame.
- **Passive Dynamics:** Perhaps the most promising approach utilizes passive dynamics where the momentum of swinging limbs is used for greater efficiency. It has been shown that totally unpowered humanoid mechanisms can walk down a gentle slope, using only gravity to propel themselves. Using this technique, a robot need only supply a small amount of motor power to walk along a flat surface or a little more to walk up a hill. This technique promises to make walking robots at least ten times more efficient than ZMP walkers, like ASIMO.

Other methods of locomotion

RQ-4 Global Hawk unmanned aerial vehicle

- **Flying:** A modern passenger airliner is essentially a flying robot, with two humans to manage it. The autopilot can control the plane for each stage of the journey, including takeoff, normal flight, and even landing. Other flying robots are uninhabited, and are known as unmanned aerial vehicles (UAVs). They can be smaller and lighter without a human pilot onboard, and fly into dangerous territory for military surveillance missions. Some can even fire on targets under command. UAVs are also being developed which can fire on targets automatically, without the need for a command from a human. Other flying robots include cruise missiles, the Entomopter, and the Epson micro helicopter robot. Robots such as the Air Penguin, Air Ray, and Air Jelly have lighter-than-air bodies, propelled by paddles, and guided by sonar.

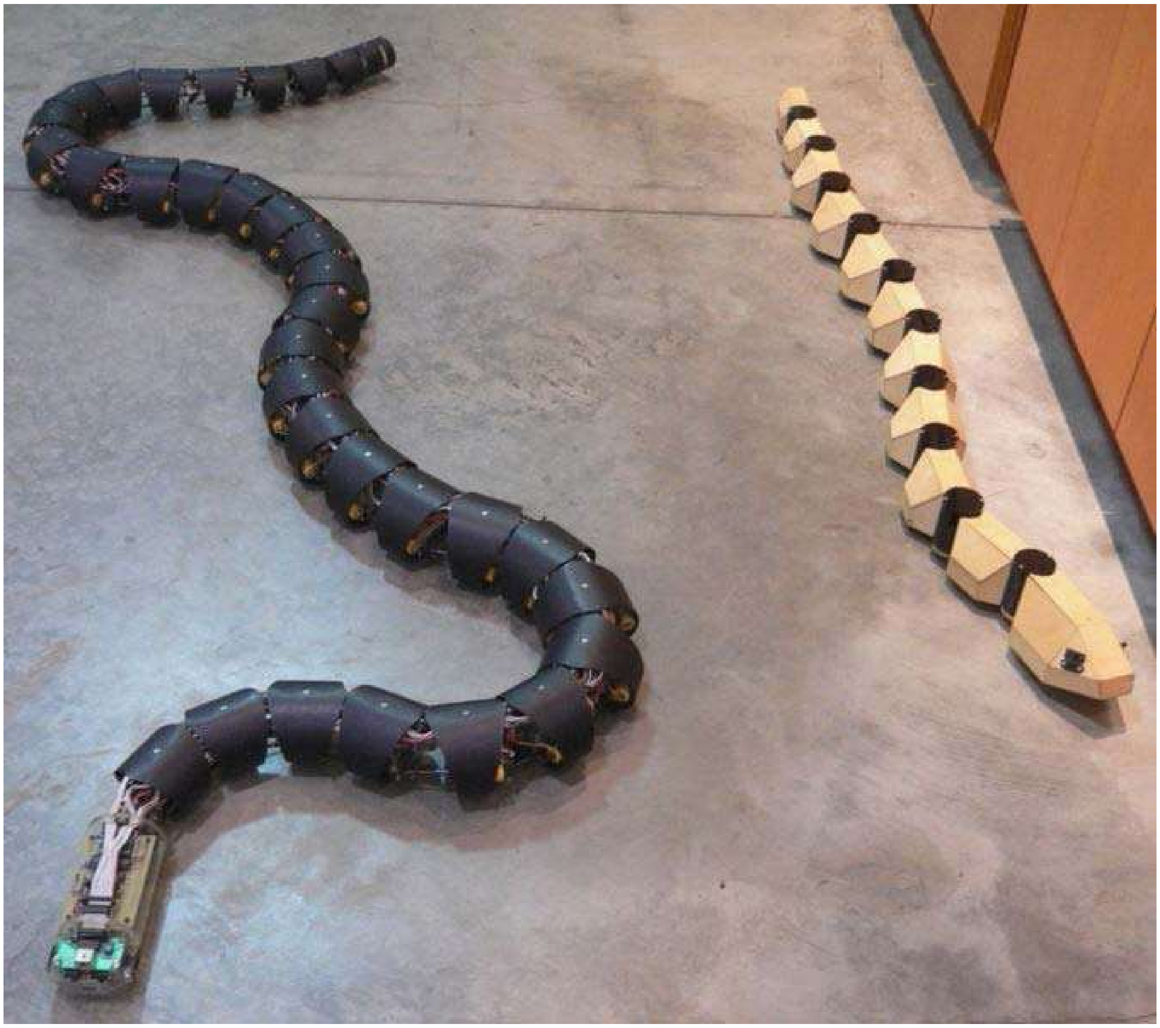

Two robot snakes. Left one has 64 motors (with 2 degrees of freedom per segment), the right one 10.

- **Snaking:** Several snake robots have been successfully developed. Mimicking the way real snakes move, these robots can navigate very confined spaces, meaning they may one day be used to search for people trapped in collapsed buildings. The Japanese ACM-R5 snake robot can even navigate both on land and in water.
- **Skating:** A small number of skating robots have been developed, one of which is a multi-mode walking and skating device, Titan VIII. It has four legs, with unpowered wheels, which can either step or roll. Another robot, Plen, can use a miniature skateboard or rollerskates, and skate across a desktop.
- **Climbing:** Several different approaches have been used to develop robots that have the ability to climb vertical surfaces. One approach mimicks the movements of a human climber on a wall with protrusions; adjusting the center of mass and moving each limb in turn to gain leverage. An example of this is Capuchin, built by Stanford University, California. Another approach uses the specialised toe pad method of wall-climbing geckoes, which can run on smooth surfaces such as vertical glass. Examples of this approach include Wallbot and Stickybot. China's

"Technology Daily" November 15, 2008 reported New Concept Aircraft (ZHUHAI) Co., Ltd. Dr. Li Hiu Yeung and his research group have recently successfully developed the bionic gecko robot "Speedy Freelander".According to Dr. Li introduction, this gecko robot can rapidly climbing up and down in a variety of building walls, ground and vertical wall fissure or walking upside down on the ceiling, it is able to adapt on smooth glass, rough or sticky dust walls as well as the various surface of metallic materials and also can automatically identify obstacles, circumvent the bypass and flexible and realistic movements. Its flexibility and speed are comparable to the natural gecko. A third approach is to mimick the motion of a snake climbing a pole .

- **Swimming:** It is calculated that when swimming some fish can achieve a propulsive efficiency greater than 90%. Furthermore, they can accelerate and maneuver far better than any man-made boat or submarine, and produce less noise and water disturbance. Therefore, many researchers studying underwater robots would like to copy this type of locomotion. Notable examples are the Essex University Computer Science Robotic Fish, and the Robot Tuna built by the Institute of Field Robotics, to analyze and mathematically model thunniform motion. The Aqua Penguin, designed and built by Festo of Germany, copies the streamlined shape and propulsion by front "flippers" of penguins. Festo have also built the Aqua Ray and Aqua Jelly, which emulate the locomotion of manta ray, and jellyfish, respectively.

Environmental interaction and navigation

RADAR, GPS, LIDAR, ... are all combined to provide proper navigation and obstacle avoidance

Though a significant percentage of robots in commission today are either human controlled, or operate in a static environment, there is an increasing interest in robots that can operate autonomously in a dynamic environment. These robots require some combination of navigation hardware and software in order to traverse their environment. In particular unforeseen events (e.g. people and other obstacles that are not stationary) can cause problems or collisions. Some highly advanced robots as ASIMO, EveR-1, Meinü robot have particularly good robot navigation hardware and software. Also, self-controlled cars, Ernst Dickmanns' driverless car, and the entries in the DARPA Grand Challenge, are capable of sensing the environment well and subsequently making navigational decisions based on this information. Most of these robots employ a GPS navigation device with waypoints, along with radar, sometimes combined with other sensory data such as LIDAR, video cameras, and inertial guidance systems for better navigation between waypoints.

Human-robot interaction

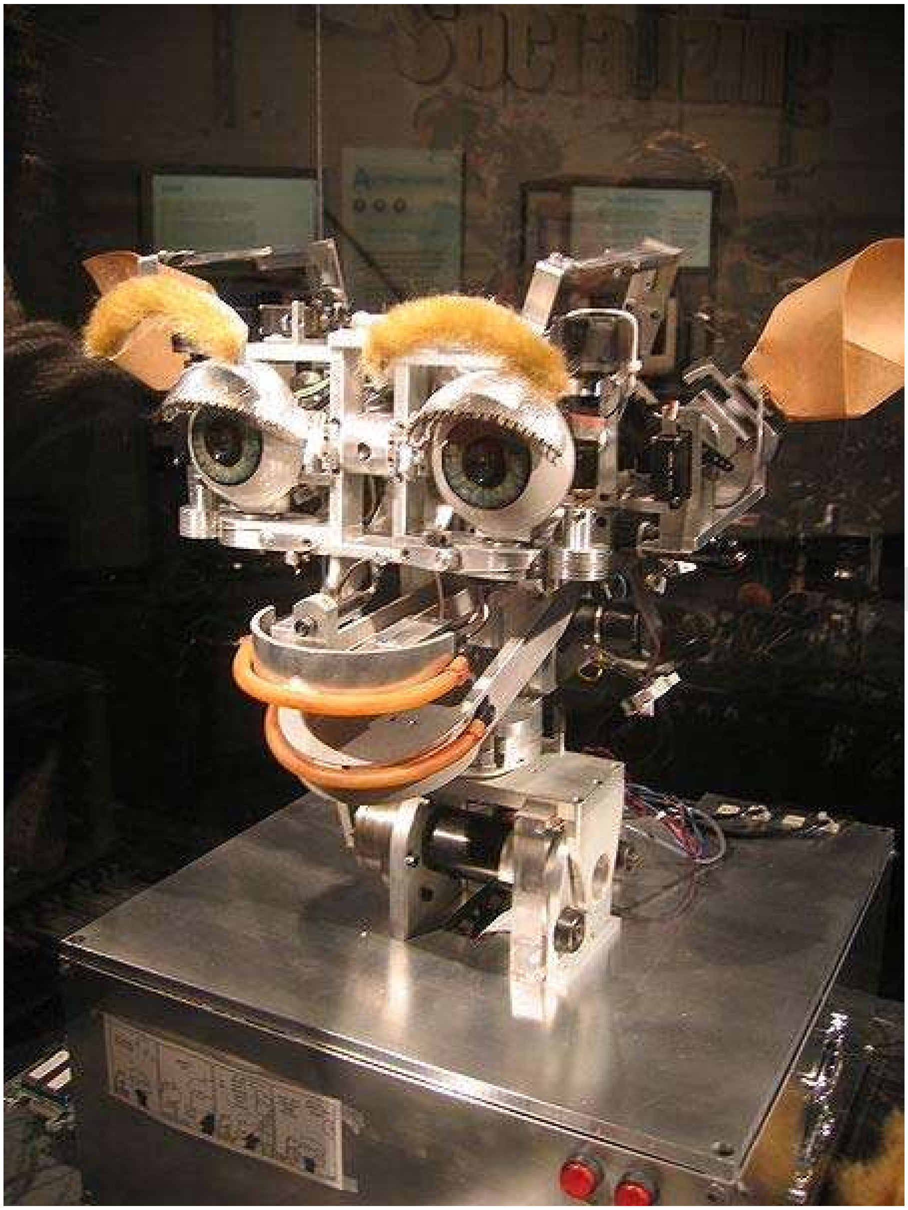

Kismet can produce a range of facial expressions.

If robots are to work effectively in homes and other non-industrial environments, the way they are instructed to perform their jobs, and especially how they will be told to stop will be of critical importance. The people who interact with them may have little or no training in robotics, and so any interface will need to be extremely intuitive. Science fiction authors also typically assume that robots will eventually be capable of communicating with humans through speech, gestures, and facial expressions, rather than a command-line interface. Although speech would be the most natural way for the human to communicate, it is unnatural for the robot. It will be a long time before robots interact as naturally as the fictional C-3PO.

- **Speech recognition:** Interpreting the continuous flow of sounds coming from a human (speech recognition), in real time, is a difficult task for a computer, mostly because of the great variability of speech. The same word, spoken by the same person may sound different depending on local acoustics, volume, the previous word, whether or not the speaker has a cold, etc.. It becomes even harder when the speaker has a different accent. Nevertheless, great strides have been made in the field since Davis, Biddulph, and Balashek designed the first "voice input system" which recognized "ten digits spoken by a single user with 100% accuracy" in 1952. Currently, the best systems can recognize continuous, natural speech, up to 160 words per minute, with an accuracy of 95%.
- **Gestures:** One can imagine, in the future, explaining to a robot chef how to make a pastry, or asking directions from a robot police officer. In both of these cases, making hand gestures would aid the verbal descriptions. In the first case, the robot would be recognizing gestures made by the human, and perhaps repeating them for confirmation. In the second case, the robot police officer would gesture to indicate "down the road, then turn right". It is likely that gestures will make up a part of the interaction between humans and robots. A great many systems have been developed to recognize human hand gestures.
- **Facial expression:** Facial expressions can provide rapid feedback on the progress of a dialog between two humans, and soon it may be able to do the same for humans and robots. Robotic faces have been constructed by Hanson Robotics using their elastic polymer called Frubber, allowing a great amount of facial expressions due to the elasticity of the rubber facial coating and imbedded subsurface motors (servos) to produce the facial expressions. The coating and servos are built on a metal skull. A robot should know how to approach a human, judging by their facial expression and body language. Whether the person is happy, frightened, or crazy-looking affects the type of interaction expected of the robot. Likewise, robots like Kismet and the more recent addition, Nexi can produce a range of facial expressions, allowing it to have meaningful social exchanges with humans.
- **Artificial emotions** Artificial emotions can also be imbedded and are composed of a sequence of facial expressions and/or gestures. As can be seen from the movie Final Fantasy: The Spirits Within, the programming of these artificial emotions is complex and requires a great amount of human observation. To simplify this programming in the movie, presets were created together with a special software program. This decreased the amount of time needed to make the film. These presets could possibly be transferred for use in real-life robots.
- **Personality:** Many of the robots of science fiction have a personality, something which may or may not be desirable in the commercial robots of the future. Nevertheless, researchers are trying to create robots which appear to have a personality: i.e. they use sounds, facial expressions, and body language to try to convey an internal state, which may be joy, sadness, or fear. One commercial example is Pleo, a toy robot dinosaur, which can exhibit several apparent emotions.

Control

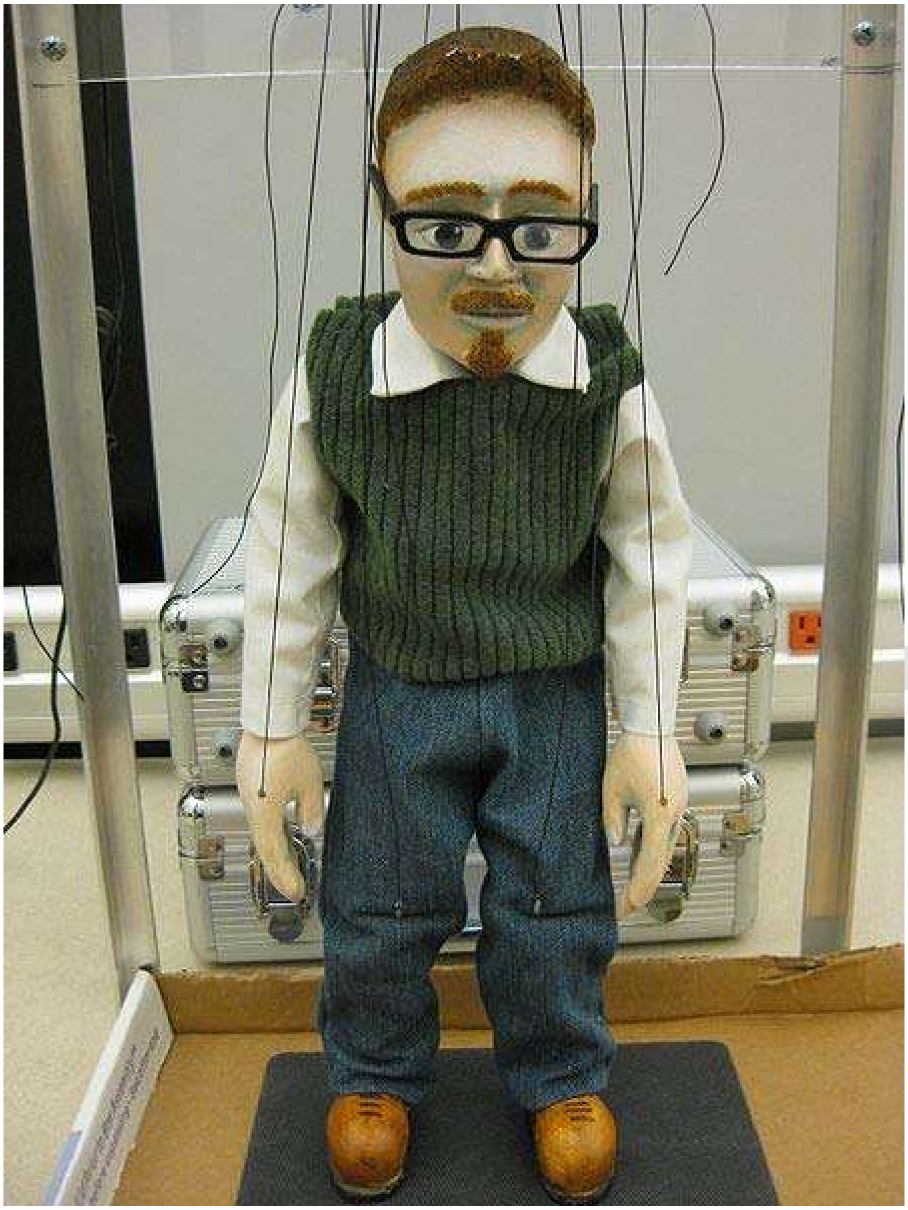

A robot-manipulated marionette, with complex control systems

The mechanical structure of a robot must be controlled to perform tasks. The control of a robot involves three distinct phases - perception, processing, and action (robotic paradigms). Sensors give information about the environment or the robot itself (e.g. the position of its joints or its end effector). This information is then processed to calculate the appropriate signals to the actuators (motors) which move the mechanical.

The processing phase can range in complexity. At a reactive level, it may translate raw sensor information directly into actuator commands. Sensor fusion may first be used to estimate parameters of interest (e.g. the position of the robot's gripper) from noisy sensor data. An immediate task (such as moving the gripper in a certain direction) is inferred from these estimates. Techniques from control theory convert the task into commands that drive the actuators.

At longer time scales or with more sophisticated tasks, the robot may need to build and reason with a "cognitive" model. Cognitive models try to represent the robot, the world, and how they interact. Pattern recognition and computer vision can be used to track objects. Mapping techniques can be used to build maps of the world. Finally, motion planning and other artificial intelligence techniques may be used to figure out how to act. For example, a planner may figure out how to achieve a task without hitting obstacles, falling over, etc.

Autonomy levels

Control systems may also have varying levels of autonomy.

1. Direct interaction is used for haptic or tele-operated devices, and the human has nearly complete control over the robot's motion.
2. Operator-assist modes have the operator commanding medium-to-high-level tasks, with the robot automatically figuring out how to achieve them.
3. An autonomous robot may go for extended periods of time without human interaction. Higher levels of autonomy do not necessarily require more complex cognitive capabilities. For example, robots in assembly plants are completely autonomous, but operate in a fixed pattern.

Another classification takes into account the interaction between human control and the machine motions.

1. Teleoperation. A human controls each movement, each machine actuator change is specified by the operator.
2. Supervisory. A human specifies general moves or position changes and the machine decides specific movements of its actuators.
3. Task-level autonomy. The operator specifies only the task and the robot manages itself to complete it.
4. Full autonomy. The machine will create and complete all its tasks without human interaction.

Dynamics and kinematics

The study of motion can be divided into kinematics and dynamics. Direct kinematics refers to the calculation of end effector position, orientation, velocity, and acceleration when the corresponding joint values are known. Inverse kinematics refers to the opposite case in which required joint values are calculated for given end effector values, as done in path planning. Some special aspects of kinematics include handling of redundancy (different possibilities of performing the same movement), collision avoidance, and singularity avoidance. Once all relevant positions, velocities, and accelerations have been calculated using kinematics, methods from the field of dynamics are used to study the effect of forces upon these movements. Direct dynamics refers to the calculation of accelerations in the robot once the applied forces are known. Direct dynamics is used in computer simulations of the robot. Inverse dynamics refers to the calculation of the actuator forces necessary to create a prescribed end effector acceleration. This information can be used to improve the control algorithms of a robot.

In each area mentioned above, researchers strive to develop new concepts and strategies, improve existing ones, and improve the interaction between these areas. To do this, criteria for "optimal" performance and ways to optimize design, structure, and control of robots must be developed and implemented.

Robot research

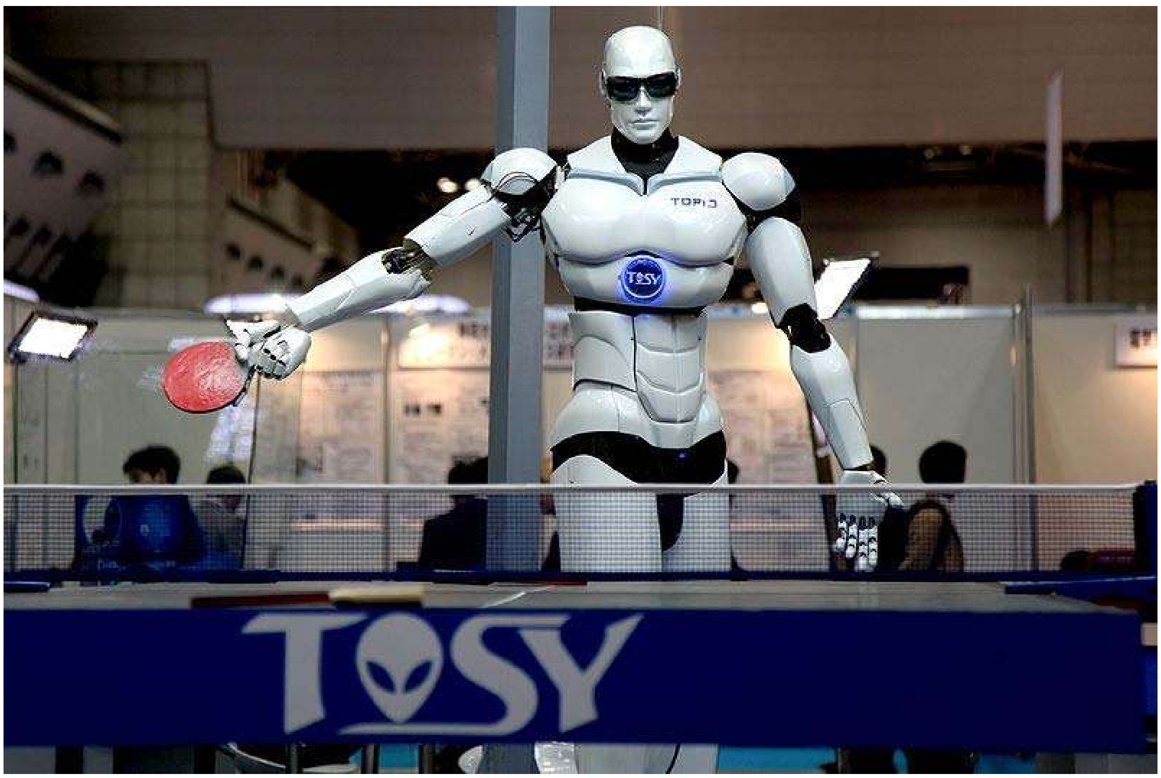

TOPIO, a robot developed by TOSY that can play ping-pong.

Much of the research in robotics focuses not on specific industrial tasks, but on investigations into new types of robots, alternative ways to think about or design robots, and new ways to manufacture them but other investigations, such as MIT's cyberflora project, are almost wholly academic.

A first particular new innovation in robot design is the opensourcing of robot-projects. To describe the level of advancement of a robot, the term "Generation Robots" can be used. This term is coined by Professor Hans Moravec, Principal Research Scientist at the Carnegie Mellon University Robotics Institute in describing the near future evolution of robot technology. *First generation* robots, Moravec predicted in 1997, should have an intellectual capacity comparable to perhaps a lizard and should become available by 2010. Because the *first generation* robot would be incapable of learning, however, Moravec predicts that the *second generation* robot would be an improvement over the *first* and become available by 2020, with an intelligence maybe comparable to that of a mouse. The *third generation* robot should have an intelligence comparable to that of a monkey. Though *fourth generation* robots, robots with human intelligence, professor Moravec predicts, would become possible, he does not predict this happening before around 2040 or 2050.

The second is Evolutionary Robots. This is a methodology that uses evolutionary computation to help design robots, especially the body form, or motion and behavior controllers. In a similar way to natural evolution, a large population of robots is allowed to compete in some way, or their ability to perform a task is measured using a fitness function. Those that perform worst are removed from the population, and replaced by a new set, which have new behaviors based on those of the winners. Over time the population improves, and eventually a satisfactory robot may appear. This happens without any direct programming of the robots by the researchers. Researchers use this method both to create better robots, and to explore the nature of evolution. Because the process often requires many generations of robots to be simulated, this technique may be run entirely or mostly in simulation, then tested on real robots once the evolved algorithms are good enough. Currently, there are about 1 million industrial robots toiling around the world, and Japan is the top country having high density of utilizing robots in its manufacturing industry.

Education and training

The SCORBOT-ER 4u - educational robot.

Robots recently became a popular tool in raising interests in computing for middle and high school students. First year computer science courses at several universities were developed which involves the programming of a robot instead of the traditional software engineering based coursework.

Career training

Universities offer Bachelors and Masters degrees in the field of robotics. Select Private Career Colleges and vocational schools offer robotics training to train individuals towards being job ready and employable in the emerging robotics industry.

Certification

The Robotics Certification Standards Alliance (RCSA) is an international robotics certification authority who confers various industry and educational related robotics certifications.

Employment in robotics

A robot technician builds small all-terrain robots.

Robotics is an essential component in any modern manufacturing environment. As factories increase their use of robots, the number of robotics related jobs grow and have been observed to be on a steady rise.

Relationship to unemployment

Some analysts, such as Martin Ford, argue that robots and other forms of automation will ultimately result in significant unemployment as machines begin to match and exceed the capability of workers to perform most jobs. At present the negative impact is only on average and repetitive jobs, and there is actually a positive impact on the number of jobs for highly skilled technicians, engineers, and knowledge workers. However, these highly skilled jobs are not sufficient in number to offset the greater decrease in employment among the general population, causing structural unemployment in which overall (net) unemployment rises.

As robotics and artificial intelligence develop further, some worry even many skilled jobs may be threatened. In conventional economic theory, this should cause merely an increase in the productivity of the involved industries, resulting in higher demand for other goods, and hence higher labour demand in these sectors, offsetting whatever negatives are caused. However, some authors believe that the conventional theory describes the past well but may not describe the future because of shifts in the parameter values that shape the context.

Healthcare

Script Pro manufactures a robot designed to help pharmacies fill prescriptions that consist of oral solids or medications in pill form. The pharmacist or pharmacy technician enters the prescription information into its information system. The system, upon determining whether or not the drug is in the robot, will send the information to the robot for filling. The robot has 3 different size vials to fill determined by the size of the pill. The robot technician, user, or pharmacist determines the needed size of the vial based on the tablet when the robot is stocked. Once the vial is filled it is brought up to a conveyor belt that delivers it to a holder that spins the vial and attaches the patient label. Afterwards it is set on another conveyor that delivers the patient's medication vial to a slot labeled with the patient's name on an LED read out. The pharmacist or technician then checks the contents of the vial to ensure it's the correct drug for the correct patient and then seals the vials and sends it out front to be picked up. The robot is a very time efficient device that the pharmacy depends on to fill prescriptions.

McKesson's Robot RX is another healthcare robotics product that helps pharmacies dispense thousands of medications daily with little or no errors. The robot can be ten feet wide and thirty feet long and can hold hundreds of different kinds of medications and thousands of doses. The pharmacy saves many resources like staff members that are otherwise unavailable in a resource scarce industry. It uses an electromechanical head coupled with a pneumatic system to capture each dose and deliver it to its either stocked or dispensed location. The head moves along a single axis while it rotates 180 degrees to pull the medications. During this process it uses barcode technology to verify its pulling the correct drug. It then delivers the drug to a patient specific bin on a conveyor belt. Once the bin is filled with all of the drugs that a particular patient needs and that the robot stocks, the bin is then released and returned out on the conveyor belt to a technician waiting to load it into a cart for delivery to the floor.