William L. WHITTAKER
Behnam MOTAZED
Construction Robotics Laboratory
The Robotics Institute
Carnegie-Mellon University

Evolution of a Robotic Excavator

Abstract

This paper considers the opportunities for general robotic excavation based on available experience with utility pipe excavation, and on the broader base of robotic technology. We first critique our robotic pipe excavator to identify shortcomings and specializations that would be limiting to more robust and general excavation systems. We then consider developments that could meet the technological needs. Finally, we present a speculative version of a generic excavation robot.

Introduction

Excavation is an excellent application to further the evolution of construction robotics because of its significance in scale and economic importance. Excavation operates on a universal and generic material (soil), and its goal and state can be adequately described by robot models of geometry and kinetics. Excavation is a model example of a class of construction tasks that defy preplanning and proceed in the exploratory mode of discovery. Further, excavation is tolerant of imprecision, well understood as a human-driven process, and prototypical of a host of spin-off applications.

Robotic excavation initiatives are emerging in diverse applications such as offshore ocean floor construction, pipe excavation for gas utilities, and crater repair of bomb-damaged runways. The need is to envision and evolve excavation robots that respond to these initiatives.

In this paper, we speculate about a cognitive robot system for General Robotic Excavation (GENEREX). We consider developments that could overcome the shortcomings and specializations of our original system, leading to a more general and robust robotic excavation.

REX

Earlier, we prototyped the first Robotic Excavator (REX) which integrated sensing, modeling, planning, simulation, and action specifically to unearth buried utility piping. REX reduced the excavation hazard posed by explosive gasses, decreased operation costs, and increased productivity in the gas utility industry.

REX's operating procedures were analogous to an archeological dig for historic remains. Soil was gently cleared from buried objects that were gradually discovered as the excavation progressed. This loop of sense-discover-excavate continued until pipes were fully exposed or until a preset depth was reached without discovering a pipe.

REX's sonar sensor collected range data to instantiate a terrain model of the excavation site. The terrain surface model was interpreted by applying edge detection and line fitting algorithms to discover features, and in turn, to recognize pipes and non-erodable objects such as rocks or buried foundations. REX then consulted this sensor-built terrain model to plan digging operations that eroded the terrain away in layers. Excavation plans were predicated on the terrain and object models that were updated after the erosion of each layer. Higher-level planning could alter the course of excavation to specifically unearth a discovered pipe.

Conventional strategies and traditional digging tools were not incorporated in REX because of the great difficulty of sensing, modeling and planning forceful tool motions automatically. The benign excavation tooling developed for REX was a supersonic air-jet cutter that dislodged material without the direct contact intrinsic to bucket excavation. This non-contact tooling considerably reduced the potential for the system to damage itself and its environment and greatly simplified autonomous tool planning. A flexible, instrumented tool holder provided the necessary compliance to make collision a gradual event and, as a last resort, bump sensing detected unforeseen collisions.

Human interfaces to REX took the forms of joystick, keyboard and animated display. The REX actuation hardware consisted of a rugged hydraulic arm appended to a four-link backhoe (modified for computer servo-control), which in turn was mounted to a utility truck. Together, the backhoe and arm combined to deploy the air-jet tooling and the sonar range sensor throughout the excavation site.

REX Evaluation

REX was a research implementation for autonomous excavation without excessive pretenses to robustness and depth. Our breadth-first thinking invested early to identify critical unknowns of great value, seeking to provide a solution to the problem, the solution itself becoming a throwaway. Our laboratory implementation of autonomous excavation returned handsomely on its investment by way of lessons learned. We reflect on those lessons here.

A positive feature of REX was that it was a complete, autonomous robot implementation, though admittedly shallow. Further, the simplicity of elements enabled quasi-real-time system performance. The challenge of meeting the real-time constraint forced a three-level planning paradigm of strategy, tactic and reflex that is now fundamental to autonomous planning and control. REX was an object lesson for insights into sensor driven manipulation and mono-modal, autonomous robot systems in general.

REX had five major liabilities. The first was weak sensors. Confused by noise and flying debris, acoustic range sensors were not sufficiently rugged for the environment and were incompatible with REX's air-jet tooling. They limited the system to a single mode of spatial sensing, modeling and planning when full cognition for excavation would require additional sensing modes. REX's second liability was weak models; simple depth maps cannot model the complex shapes created by general excavation. The third was weak planning. Because REX had no cognitive ability, it could not plan excavation completely for any useful problem. Tooling, the fourth liability, was appealing for its ease of sensing, modeling and planning of non-contact processes, but was unable to achieve feasible productivity and was unable to excavate under many site conditions. And finally, REX had no force cognizance, and thus had to avoid all contact. Non-contact tooling is impotent and inappropriate for production excavation. The robots and materials of excavation are not so fragile that they are damaged by controlled forces.

GENEREX

Our research has calibrated and filled out our perspectives on REX and identified its deficiencies through a first experimental implementation. We are currently improving tooling, manipulation, control, sensing, modeling, and strategy for a second-generation excavator, GENEREX, with refinements on the original.

Excavator Architecture

Our research toward a cognitive excavator is based on a rational framework that focuses our component technology research and our integration effort. Figure 1 shows the framework of the proposed excavator robot, in which components at various levels fit together to form an integrated system.

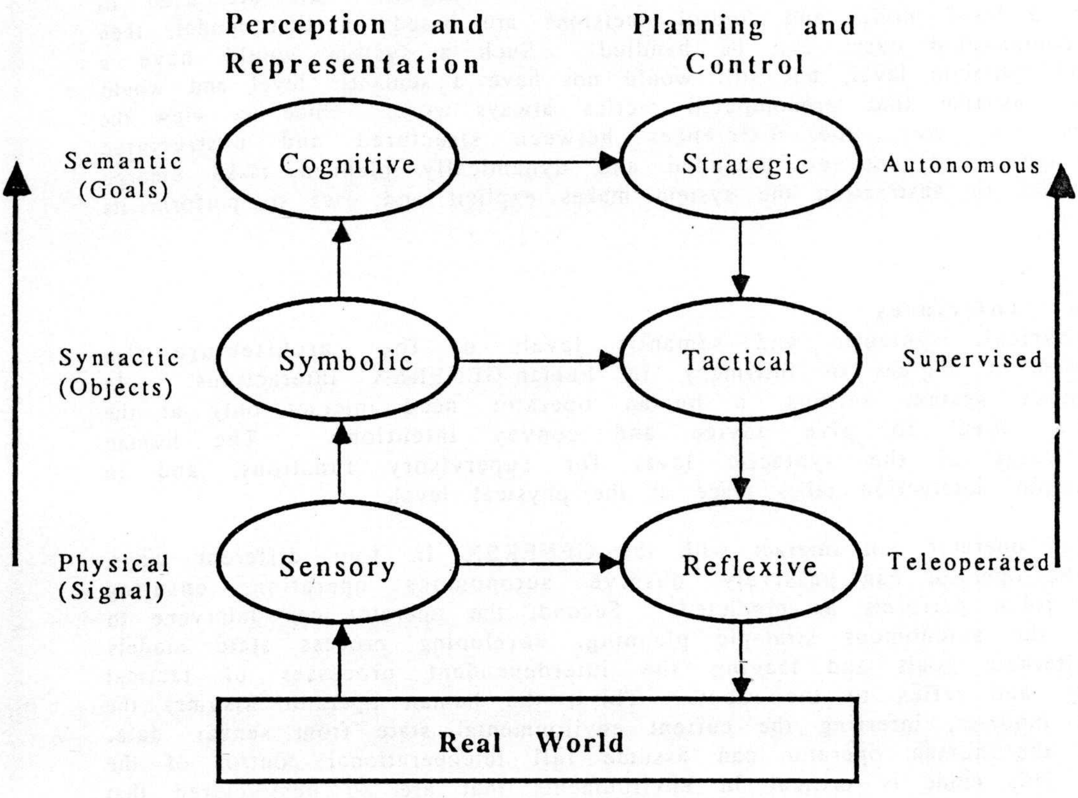

Figure 1: GENEREX System Architecture

GENEREX uses multiple sensors to acquire raw visual, tactile, force, and magnetic sensory data. The result of interpreting sensory data is represented in a symbolic manner accessible by modules for cognition and tactical planning. The highest level involves cognition, planning, and learning by relating the symbolic representation of scene and object with the goals, missions, and purposes of GENEREX. At the lowest level of action, the system actuators are driven either reflexively by sensory data or by tactical plans. The strategic course of action is planned using cognitive interpretation of sensory data along with *a priori* world knowledge to evaluate the status and mission of the system operation. Between strategic and reflexive levels lies the tactical planning that uses the interpreted sensory information to generate sequences of executable operations.

The framework presents three distinct levels of abstraction. At the bottom lies the signal and physical level, encompassing sensory data and reflexive action. The intermediate, syntactic level represents the robot's understanding and action in symbolic form. The highest level includes semantic concepts such as "purpose," "goal," and "mission."

The capabilities active in GENEREX depend on how high a level of abstraction is explicitly understood and handled at the time. For example, when the robot is servoing on a magnetic field it has no explicit understanding of anything above the sensory and reflexive level. When magnetic data are used to constuct a local model and control decisions are based on that model, then more complicated cases can be handled. Such a system would have a complete syntactic level, but still would not have a semantic level and would therefore assume that precomputed tactics always work. Once we view the problem this way, the differences between structured and unstructured environments and between precoded and dynamically planned tasks emerge in the level of abstraction the system makes explicit and uses to perform its mission.

Human Interfaces

The physical, syntactic and semantic levels of the architecture also correspond to degrees of initimacy in human-GENEREX interactions. In autonomous system actions, a human operator need interact only at the semantic level to give advice and convey intentions. The human communicates at the syntactic level for supervisory functions, and in teleoperation, interaction takes place at the physical level.

A human operator can interact with the GENEREX in four different ways. First, the operator can passively observe autonomous operation, ensuring that the robot performs as predicted. Second, the operator can intervene to displace the autonomous strategic planning, developing process state models from alternate goals and leaving the interdependent processes of tactical planning and reflex to the robot. Third, the human operator assumes the role of modeler, inferring the current environmental state from sensor data. Finally, the human operator can assume full teleoperational control of the robot. This mode is critical in environments that are so unstructured that neither human nor robot can completely preplan against contingencies.

Spatial Cognition

Spatial cognition, the ability to model and reason about geometry, is a useful mode of construction robot planning. A surface-specific strategy first interprets a model of a terrain surface, and then servoes to that surface, executing simple tasks such as "fly over the surface," "point toward the surface," or "erode the surface by a given increment." Promising applications of surface-specific strategies include concrete form cleaning, sealing foundations, sandblasting barges, repairing airstrips, washing airplanes, dredging seabottom, spray washing, surface removal and excavating pipes, as implemented here. The reflexes and tactics for these vast and important tasks are manifested in similar hardware and software; the significant differences are mostly strategic.

Generalization of these tactics will be invaluable to the advance of perceptive robots in unstructured environments. Thus, we are exploring the advantages of reflex tactics based on terrain models. We believe that data structures for surface representation coupled with tactical servoing allow more relevant methods of organization, access, and use relative to symbolized data and geometry for certain problems in the domain. Both surface models and planning schemes can generalize or refine data for use at variable resolutions (i.e., big picture to details). Common tactics can be made simple enough to generate the fast, interactive trajectory reflexes that become the lower nervous system of goal-directed, autonomous workers. While symbolic approaches remain essential at the highest planning level for many activities, they cannot compete with the simplicity, speed, and relevance of the surface-specific tactics discussed here. We suspect that symbolic approaches may not even prove relevant to excavation, grading, and earthworks; these domains are essentially devoid of feature, which is to say, all dirt looks alike.

Voluminous research is rightly invested in image understanding and interpretation of range data to infer symbolic world models. (For example, inferential pipe modeling is critical in the higher, strategic planning of our excavation exercise.) The hope of many significant research programs is that symbolic understanding will become sufficiently fast and robust for cognitive robots to strategize their actions in the top-down manner of manufacturing robots. We believe, however, that this paradigm does not take advantage of the tremendous relevance of the instant reflexes and tactical planning that are possible from direct use of data without resorting to higher abstraction. Our current view is that reflex and subliminal tactics (millisecond gameplans) are as essential to cognitive robots as they are to higher animals.

Because known range sensing hardware is incompatible in the excavation environment, the full ideal of spatial cognizance can not be fulfilled in robotic excavation. Specifically, range sensing is unreliable during the digging cycle, so it is not a useful observable for reflexive response. However, the force cognizance of the GENEREX motion hardware offers an alternative. Motion hardware is forgiving in coollision, and so is useful as a reflexive sensing mode, allowing motion without the assurances that full spatial cognizance provides. It is possible to use GENEREX manipulator motion as a tactile agent, and to interpret the volumes swept by the hardware as void. Once integrated, this model of free space is useful for spatial planning.

Force Cognition

Kinetic control and a force-specific strategy (e.g., digging by feel in a sandbox) could revolutionize benign excavation. Damage control is akin to force control, and though we want to preclude unwanted damage, our implementation to date ignores force. Instead of planning around a kinematic ballet, we plan to implement a robotic excavator with force cognizant manipulation, a technique where digging is controlled by feel, applying sufficient, but never excessive, force. Force cognizant manipulation emulates a human with a hand shovel. While digging, a human receives a qualitative measure of force. The robotic excavator, however, has advantages over a human-operated backhoe because simple backhoe operation provides no measure of force control. Quantitative measures of force feedback to the control system thus enhance system robustness. While kinematics still play an important role, forces become the determining basis for planning and action.

GENEREX will contact and engage the soil, overcoming the limitations encountered in the non-contact scenario. These techniques are applicable to an entire class of problems in robotics where kinematic and spatial planning alone are insufficient solutions.

Magnetic Cognition

There is great promise and relevance for excavation robots that can reason about ferrous objects encountered in the ground. Buried man-made objects like pipes, cables and tracer wires are of great interest to digging robots and the humans they serve. Magnetic field sensing is the key to perceiving objects before, during and after excavation with the speed and confidence to guide the digging. This is possible at several layers of abstraction with or without the soil intervening. Models comparable to those implemented for range data can be built from instantaneous readings, field models and inferring from full fields. Tactical and reflexive plans and responses are possible because range information is readily available. Excavation robots will be able to servo directly on raw magnetic observables. These rudiments of magnetic sensing, modeling and planning have shown that fuller magnetic cognizance can be developed.

Sensors

There are no forseeable breakthroughs in range sensing that are compatible with the scream of air-jet tooling, with flying excavation debris and with the impact and vibration of excavation, so spatial models must be discounted somewhat for autonomous excavation. Alternately, manipulator force sensing, internal to an arm, is innately survivable and oblivious to noise, dust and the like, making it a desirable candidate for sensing. Magnetic sensors have no moving parts, and can be hardened against dust, moisture and impact. The fusion of multiple sensor data in GENEREX will allow fuller perception of complex environmental phenomena. GENEREX will incorporate range sensing when it is possible to do so, but will depend on force and magnetic sensing for its reflexes.

Hardware

The ideal hardware configurations are not fully apparent for robotic excavator systems. It is not our intention to emphasize tooling, manipulation and locomotion hardware concerns, but they do impact the specifics of the digging process, and hence the demands upon sensing, modeling and planning for any specific application. This is important because manipulation hardware affects envelope, solution time, collision avoidance, force tactics, dynamics, and much more. Tooling matters most of all, as its effectiveness is governed by the simplicity of the tactics and strategies that can be brought to bear.

Summary

The General Robot Excavator (GENEREX) differs from its predecessors in degree, modes and merging of cognitive reasoning. It will display new standards of capability, survivability and compatability relative to excavation environments. GENEREX will accommodate human interactions in any of the three modes of autonomy, supervision or teleoperation and cognition will be multi-modal, multi-level and more extensive than earlier systems. Spatial, force and magnetic modes are leading candidates because of the ruggedness of their hardware and the relevance of their data to reasoning about excavation.

Acknowledgements

Complementary funding for the research presented here was furnished by the Western Pennsylvania Advanced Technology Center, the Gas Research Institute, the Southern California Gas Corporation, the Dravo Corporation and Briggs Technologies.

References

[1] Carbonell, J., Kanade, T. and Whittaker, W., *Robot Work Systems*, Proposal to the Office of Navy Research, The Robotics Institute, Carnegie Mellon University, April 1986.

[2] Akdogan, A. Rasim, *Surface Following Strategy for a Robotic Excavator*, Masters Thesis, Carnegie Mellon University, 1985.

[3] Baker, C. Nelson, *Simulation of a Monitor/Controller for a Robotic Excavator*, Masters Thesis, Carnegie Mellon University, 1986.

[4] Balash, A. John, *Process Simulation for a Robotic Excavator*, Masters Thesis, Carnegie Mellon University, 1985.

[5] Skibniewski, M. and Hendrickson, C., "Evaluation Method for Robotics Implementation: Application to Concrete Form Cleaning," *Second International Conference on Robotics in Construction*, Carnegie Mellon University, June 1985.

[6] Swetz, A. Scott, *Robot Pipe Excavator: A System View*, Masters Thesis, Carnegie Mellon University, 1986.

[7] Turkiyyah, M. George, *Surface Modeling and Robot Planning from Range Data*, Masters Thesis, Carnegie Mellon University, 1986.

[8] Whittaker, W. et al., "First Results in Automated Pipe Excavation," *Second International Conference on Robotics in Construction*, Carnegie Mellon University, June 1985.

[9] Whittaker, W., "Cognitive Robots for Construction," *Annual Research Review*, The Robotics Institute, Carnegie Mellon University, 1986.

William L. Whittaker
Behnam Motazed

Construction Robotics Laboratory
The Robotics Institute
Carnegie Mellon University
Schenley Park, Pa. 15213 USA
Phone: 412-268-2952