mean that populations with clearly derived Neanderthal features existed in Europe along with other populations that have much less evidently derived Neanderthal features (for example, specimens from Aragó, France, and, more significantly, Mauer, Germany). If this were the case, there would have been a very large diversity of human populations in Europe.

Further evidence and hypotheses. Besides the appearance of Neanderthal features in Middle Pleistocene hominin populations, a major change in the archeological record is detected at around 0.6–0.5 Mya. Prior to this date, lithic (stone) industries appearing in the sites are of mode I or Oldowan technology (simple flaked cobbles or blocks of stone). However, after this date, mode II lithic tools or Acheulean technology (using bifacial knapping, that is, breaking off pieces to shape a tool so that both sides are flattened to form a V-shaped cutting edge), which first originated in Africa, start to appear in the European archeological record.

The cultural changes observed in the Paleolithic record of Europe can be interpreted as the result of cultural diffusion from Africa, but with genetic continuity of local populations. An alternative hypothesis holds that *H. antecessor* was genetically replaced (or absorbed) by a new wave of immigrants coming from Africa. These new colonizers were the bearers of the new mode II technology, and they were also the direct ancestors of the Neanderthals. Currently, both possibilities, continuity and replacement, can be defended. A major argument against local continuity is the lack of derived Neanderthal features in the TD6 hominins, although some dental features may support the hypothesis of genetic continuity. In addition, a common morphological background could be put forward as evidence for the possible genetic links between Early and Middle Pleistocene populations.

Thus, a major question is whether *H. antecessor* populations evolved locally in Europe to give rise to the ancestors of the Neanderthals, for example, *H. heidelbergensis*. The issue of continuity could also be applied to the link between TE9 hominins (1.2–1.1 Mya) and TD6 hominins (0.8 Mya). Did they belong to the same wave of colonizers to Europe? The possibility of different waves/species arriving in Europe is theoretically possible. Given the present evidence, the most parsimonious hypothesis is that both samples came from the same ancestral species.

These various questions are linked by the issue of genetic relationships between human populations throughout the Pleistocene of Europe. The null hypothesis to be tested in the coming years, which is directly derived from the Atapuerca finds, is whether or not there is continuity of the lineages represented by *H. antecessor*, *H. heidelbergensis*, and *H. neanderthalensis*.

For background information *see* ANTHROPOLOGY; ARCHEOLOGY; DATING METHODS; EARLY MODERN HUMANS; FOSSIL HUMANS; NEANDERTALS; PALEOLITHIC; PALEONTOLOGY; PHYSICAL ANTHROPOLOGY; PLEISTOCENE; PREHISTORIC TECHNOLOGY in the McGraw-Hill Encyclopedia of Science & Technology.

Antonio Rosas

Bibliography. E. Aguirre et al., The Atapuerca sites and the Ibeas hominids, *Hum. Evol.*, 5:55–57, 1990; J. L. Arsuaga, *The Neanderthal's Necklace: In Search of the First Thinkers*, John Wiley, Chichester, U.K., 2003; J. M. Bermúdez de Castro et al., The Atapuerca sites and their contribution to the knowledge of human evolution in Europe, *Evol. Anthropol.*, 13:25–41, 2004; E. Carbonell et al., Out of Africa: The dispersal of the earliest technical systems reconsidered, *J. Anthropol. Archaeol.*, 18:119–136, 1999; A. Rosas, Human evolution in the last million years: The Atapuerca evidence, *Acta Anthropolog. Sinica Suppl.*, 19:47–56, 2000.

Autonomous passenger vehicles

Today, most industrial robots operate while bolted to factory floors, restricted to operating within steel cages in order to protect people from injury. Conveyor belts bring raw materials within reach of each robot for assembly, welding, inspection, and so forth, and take the finished products away. Such settings are called structured environments. Extraordinary human effort is required to prepare the environment, and design the manufacturing processes within that environment, so that robots can function effectively and efficiently there.

A new stage has begun in the development of robotics, in which mobile robots are beginning to leave the factory to move and work alongside people in environments designed for humans rather than for robots. Modern robots have begun to walk, roll, fly, and swim anywhere people can go—even in places where it is too dangerous for people to go, such as inside active volcanoes or near explosives. As robots migrate into environments historically occupied by people, they must become capable of operating without the crutch of a tailor-made world.

One promising research area has been the development of autonomous, or self-driving, vehicles. The goal of autonomous vehicle research is the realization of full-size passenger vehicles that can drive with no human involvement beyond the specification of a destination. One can think of such vehicles as providing the same service as a taxi, but without the human taxi driver. The key technical challenges inherent in robotic driving include handling roadways that have not been encountered before, and reacting appropriately to unpredictable vehicles, pedestrians, and other aspects of the environment.

Why develop self-driving cars? Widely available self-driving cars would accrue a number of immediately apparent societal benefits, including increased safety, productivity, and energy efficiency.

Car accidents cause more than 40,000 fatalities annually in the United States alone and about 1.3 million worldwide. Self-driving cars would reduce the frequency and severity of accidents, thereby reducing the attendant human misery and societal costs of these events.

Commuters devote hours of intense attention to the task of surviving the drive from home to work and back each day. Self-driving cars would free former drivers to use their brains for other purposes, hugely increasing productivity.

Self-driving cars, even those acting independently of other vehicles and with no direct influence on the traffic infrastructure, can drive smoothly, saving gasoline and decreasing the amount of energy dissipated as useless heat by braking. Self-driving cars capable of cooperating with other vehicles, or with the signaling infrastructure, could save even more energy, for example, by eliminating stop-and-start driving and reducing waiting at traffic lights. Another potential benefit of self-driving cars would be the elimination of the asymmetry in the burdens borne by carpool drivers and passengers.

Finally, for researchers, the development of self-driving cars promises to provide decades of intellectual excitement and engineering challenges.

What is required for autonomous driving? Seven principal technical challenges must be met in order to construct a safe, self-driving car capable of traveling along existing road networks and interoperating with existing human-driven traffic: drive-by-wire operation, representation, perception, planning, control, the human-robot interface, and the creation of a suitable platform.

Drive-by-wire operation. Any robot vehicle is based on a "drive-by-wire" mobility platform, in which the car's physically controllable "degrees of freedom"—including its transmission, gas, brake, and steering—must be placed under the control of a computer, typically through the use of electric motors and suitable hardware, firmware, and software to control them.

Representation. A robot car must have a useful internal representation of roadways (how they are delineated, by curbs, grass, painted markings, and so forth; and how they start, stop, split, and merge) and of traffic (how to safely start, follow, merge, turn into or across traffic, and stop). The robot must also have valid predictive models of how other entities in the world—trucks, cars, motorcycles, and pedestrians—are likely to behave at any moment. This behavior is of course highly dependent on the entity's location and surroundings: travel lane, exit or entrance lane, sidewalk, crosswalk, and so forth.

Perception. Of course, no vehicle can make useful progress in the world without an accurate, timely model of its surroundings. A robot car must continuously answer, and update its answers to, questions such as: Where is the drivable road surface? Are there hazards (such as potholes or road debris), obstacles (such as traffic islands, curbs, bollards, or other vehicles), or vulnerable entities (such as motorcyclists, bicyclists, or pedestrians) in or near the space likely to be occupied soon by the car? If so, how are these entities likely to move or change as the car approaches?

In practice, answering such questions involves selecting suitable sensors, positioning them in or on the vehicle so as to provide sufficient observations of the vehicle's surroundings, and integrating software to process the raw sensor data gathered as the vehicle moves into a useful interpretation of what entities are close to the vehicle or likely to come close in the near future. Sensor integration requires spatiotemporal calibration of the sensors, that is, knowledge of how data from different sensors is related across space and time. Detection and classification of objects near the vehicle is highly challenging, in part because the sensors and interpretation algorithms available today are far less sophisticated than the human visual and cognitive systems.

Planning. Once the car has an underlying representation of the world, and can perceive its surroundings, it must be able to make progress toward its "goal," that is, the destination specified by its user. (The user may or may not be a passenger; for example, a car could be summoned to its user from a parking area, or dispatched by its user to deliver a package.) In robotics terminology, the car must "plan" actions that advance it toward its goal, if possible.

Fig. 1. **Cargo compartment of Team MIT's vehicle at the 2007 DARPA Urban Challenge, which carried more than 40 central processing units (CPUs) in a built-in air-cooled mobile machine room, illustrating the enormous computational resources (by current standards) required for autonomous driving. About half of these CPUs were dedicated to machine vision processing used for finding painted lane markings on the pavement. (*Jason Dorfman, CSAIL/MIT*)**

Fig. 2. **Team MIT's vehicle, "Talos," driving unoccupied at the 2007 DARPA Urban Challenge. (*Jason Dorfman, CSAIL/MIT*)**

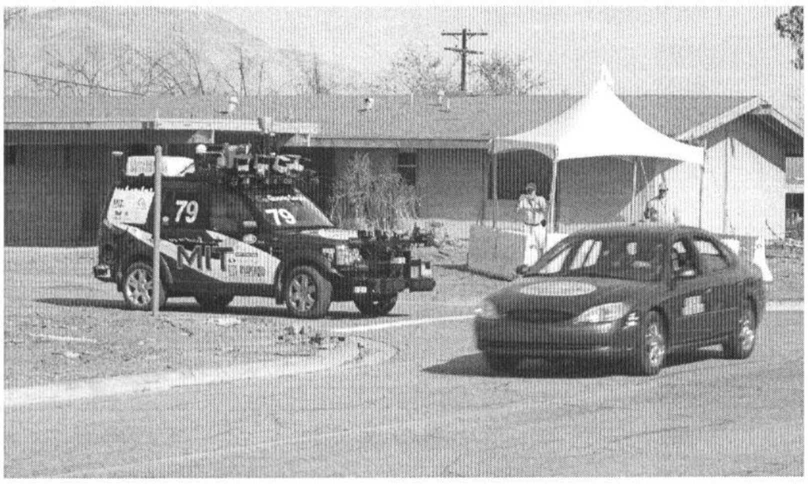

Fig. 3. Intersection on the 2007 DARPA Urban Challenge course, with human observers behind a concrete barrier. Because of the unpredictability of the current generation of robotic vehicles, DARPA removed all pedestrians from the Urban Challenge course and arranged for observers to remain behind barriers for safety. (*Jason Dorfman, CSAIL/MIT*)

Such planning involves high-level tasks, such as localizing and orienting within the world (for example, choosing to head north); medium-level tasks, such as choosing a particular north-bound road or lane; and low-level tasks, such as deciding to divert briefly around a pothole, parked car, or other vehicle. The planning task has a static aspect, in which prior information about road topology is used, and a dynamic aspect, in which live sensor information (for example, about debris on the road) or infrastructural information (for example, about traffic conditions ahead) is used. Finally, in case of unexpected occurrences such as road blockages, the vehicle must be able to replan, that is, abandon its current low-level, medium-level, or high-level plan and generate a new one. And, of course, all robots must be able to gracefully handle the case in which no feasible plan exists, that is, in which no progress may be made toward the goal. In the case of a robot vehicle, the proper behavior in this circumstance might be to wait patiently, or request help.

Control. At any moment, an autonomous vehicle must maintain a low-level plan consisting of a spatiotemporal trajectory through the world that will advance it toward its medium-level goal. For instance, the low-level plan might include a curving path through an intersection, chosen so as to exit the intersection roughly centered and aligned with the outgoing lane. Such plans also include desired bounds on speed, in addition to those for position and orientation.

The "control" problem in robotics consists of exercising the vehicle's drive-by-wire degrees of freedom so as to achieve a desired trajectory. To achieve robust control, the vehicle's "dynamics"—its response to longitudinal and lateral forces, and to gravity and road conditions—must be characterized, typically through a calibration procedure undertaken before mission-critical operation is attempted. Analogous to replanning, the vehicle's control objectives may have to be reconsidered on very short time scales in the event of unexpected occurrences such as encountering wet pavement or having a tire blow out.

Interface. The sixth essential element of a useful robotic vehicle capability is the robot's interface, that is, the means by which its human user tells the robot what to do (for example, through text or speech).

Platform. Designers of autonomous driving systems must also keep in mind the critical engineering constraint that any proposed implementation of the elements listed above will require significant resources, including computation, memory, storage, and networking, as well as internal and external sensors. Any plausible solution must provide sufficient resources for correct operation, while observing operational limits on power, volume, weight, temperature, and so forth (**Fig. 1**).

DARPA Urban Challenge. The state of the art in robotic vehicle development is well represented by the group of vehicles competing in the final round of the Defense Advanced Research Projects Agency (DARPA) Urban Challenge in November 2007 (**Fig. 2**). DARPA provided a challenge course somewhat more structured than a car would face in the fully general driving task, by clearing pedestrians from the course (**Fig. 3**), and employing only intersections involving stop signs. Moreover, DARPA supplied detailed prior information about the course to each competing vehicle, in the form of a USB stick with a plaintext Road Network Description File (RNDF) containing highly accurate Global Positioning System (GPS) coordinates for "waypoints" along each roadway, the topological and geometric properties of each intersection, and the location of all stop signs. There were no traffic signals, and all informational signs such as speed limit signs were encoded in the same plaintext file, so that the robots did not need to perceive signage in the environment.

All this information was provided to each robot 48 hours before the start of the competition, and each team was allowed to manually edit and annotate

Fig. 4. One of the world's first accidents between robotic vehicles, involving vehicles from Cornell and MIT. Each vehicle exhibited nonhuman driving behavior that was later found to have contributed to the accident. (*Jason Dorfman, CSAIL/MIT*)

the roadmap in order to provide denser waypoint information to its vehicle. Finally, after one vehicle was immobilized by its failure to achieve high-quality GPS reception, DARPA took steps to remove sources of GPS interference from the course.

In short, DARPA took many steps to make things easier for the robots, but even under these conditions the competition was dauntingly difficult. DARPA arranged for dozens of human-driven cars to move through the course. In addition, the fact that all robots would be operating simultaneously, yet independently, introduced significant randomness into the competition: A robot might predict the behavior of a "polite" human driver fairly accurately, but it proved impossible for even the human observers to predict the behavior of the robots (**Fig. 4**).

Throughout the qualifying rounds and final competition, the human-robot interface was limited to two channels: a Mission Description File (MDF) and a "remote E-stop" switch. Each MDF consisted of a list of waypoints from the RNDF, to be visited in order.

The remote E-stop switch, designed to bring a robot to a rapid halt if necessary, was controlled by the human driver of a chase car assigned to each competing robot, and by other DARPA personnel observing the course. Of 89 original entering teams, 35 were judged by DARPA as sufficiently capable to participate in the National Qualifying Event (NQE), in October 2007. Those 35 vehicles were subjected to a number of individual performance tests by DARPA, which judged 11 vehicles capable enough to participate in the final competition. Of those 11 vehicles, 5 were removed during competition, either because they drove in an unsafe manner, or simply got stuck. Of the original 89 entrants, six vehicles managed to cross the finish line.

Prospects. Over the coming decades, every aspect of robotic vehicles will improve, enabling self-driving vehicles to match, and eventually exceed, the capability of human drivers. Sensors will require less power, observe wider fields of view and longer ranges, and provide faster refresh rates. Improved radars and signal processing algorithms will enable future vehicles to "see" through fog, rain, snow, and dust. Algorithms for interpreting sensor data will become more capable, achieving (for example) more accurate classification of roads, buildings, static obstacles, vehicles of multiple types, and pedestrians. Computational resources and network bandwidth will continue to increase, and storage systems will become more capacious. Predictive models of vehicles and pedestrians will improve. Vehicles will communicate with one another to detect and avoid imminent collisions, and with the roadway infrastructure to improve traffic flow and decrease energy usage. *See* INTERVEHICLE COMMUNICATIONS.

Human-robot interfaces will improve, getting nearer to the "holy grail" of entirely natural interaction such as that between humans. Robots will understand spoken commands, and will ask questions to clarify user intent when needed. Robot vehicles will also interact with humans other than their user. For example, a robot vehicle may indicate to pedestrians that it is aware of their presence, use directional sound to warn pedestrians or other vehicles of imminent dangers, or even use visible light to "paint" areas of the road surface that will be, or may be, occupied by the vehicle in the near future.

For background information *See* COMPUTER VISION; CONTROL SYSTEMS; DRONE; GUIDANCE SYSTEMS; ROBOTICS; UNDERWATER VEHICLES in the McGraw-Hill Encyclopedia of Science & Technology.

Seth Teller

Bibliography. J. Bares and D. Wettergreen, Dante II: Technical description, results, and lessons learned, *Int. J. Robot. Res.*, 18(7):621–649, 1999; L. Fletcher et al., The MIT-Cornell collision and why it happened, Special issue on the 2007 DARPA Urban Challenge, part 3, *J. Field Robot.*, 25(10):775–807, 2008; J. Leonard et al., A perception-driven autonomous urban vehicle, Special issue on the 2007 DARPA Urban Challenge, part 3, *J. Field Robot.*, 25(10):727–774, 2008.

Avalanches and phase transitions

Many systems change from one phase to another as a function of some external parameter. For example, in boiling, liquid changes to gas as a function of temperature or pressure. The transition can occur continuously or abruptly from a starting phase, through an intermediate or a mixed phase (both phases coexisting), until reaching the final phase. In the boiling example, the transition as a function of temperature will be abrupt, with more and more liquid turning to gas at the transition temperature until there is only gas and the transition ends. In recent decades, a growing number of examples have been found of systems in which the transition behaves differently. In such systems, while the relevant parameter is changing, there are many occurrences in which a discrete "amount" of one phase changes suddenly to the other in what is called an avalanche event. One can find characteristics that are universal across different systems that have a phase transition through avalanches and do not depend on the detailed mechanism of the phase transition. The avalanches span a broad range of sizes and show a power-law distribution of avalanche magnitude; that is, there are many small avalanches, fewer medium-size ones, and only a few big avalanches. This behavior appears also in complex systems that do not possess a phase transition at all, pointing toward an even more general behavior of complex systems. Transitions through avalanches are also referred to as crackling phenomena.

Avalanches. The term avalanche is borrowed from the phenomena of snow or land avalanches. An avalanche can occur in a system that is not in equilibrium. A small perturbation to the system can have an effect that does not scale with the perturbation. For example, a skier (the perturbation) could trigger a massive amount of snow to slide down a mountain. The state before the avalanche is called a metastable state. Systems that are out of equilibrium can advance