CEILBOT

A ceiling-walking robot

Juhana Leiwo
Marco Torti
Helsinki University of Technology
2009

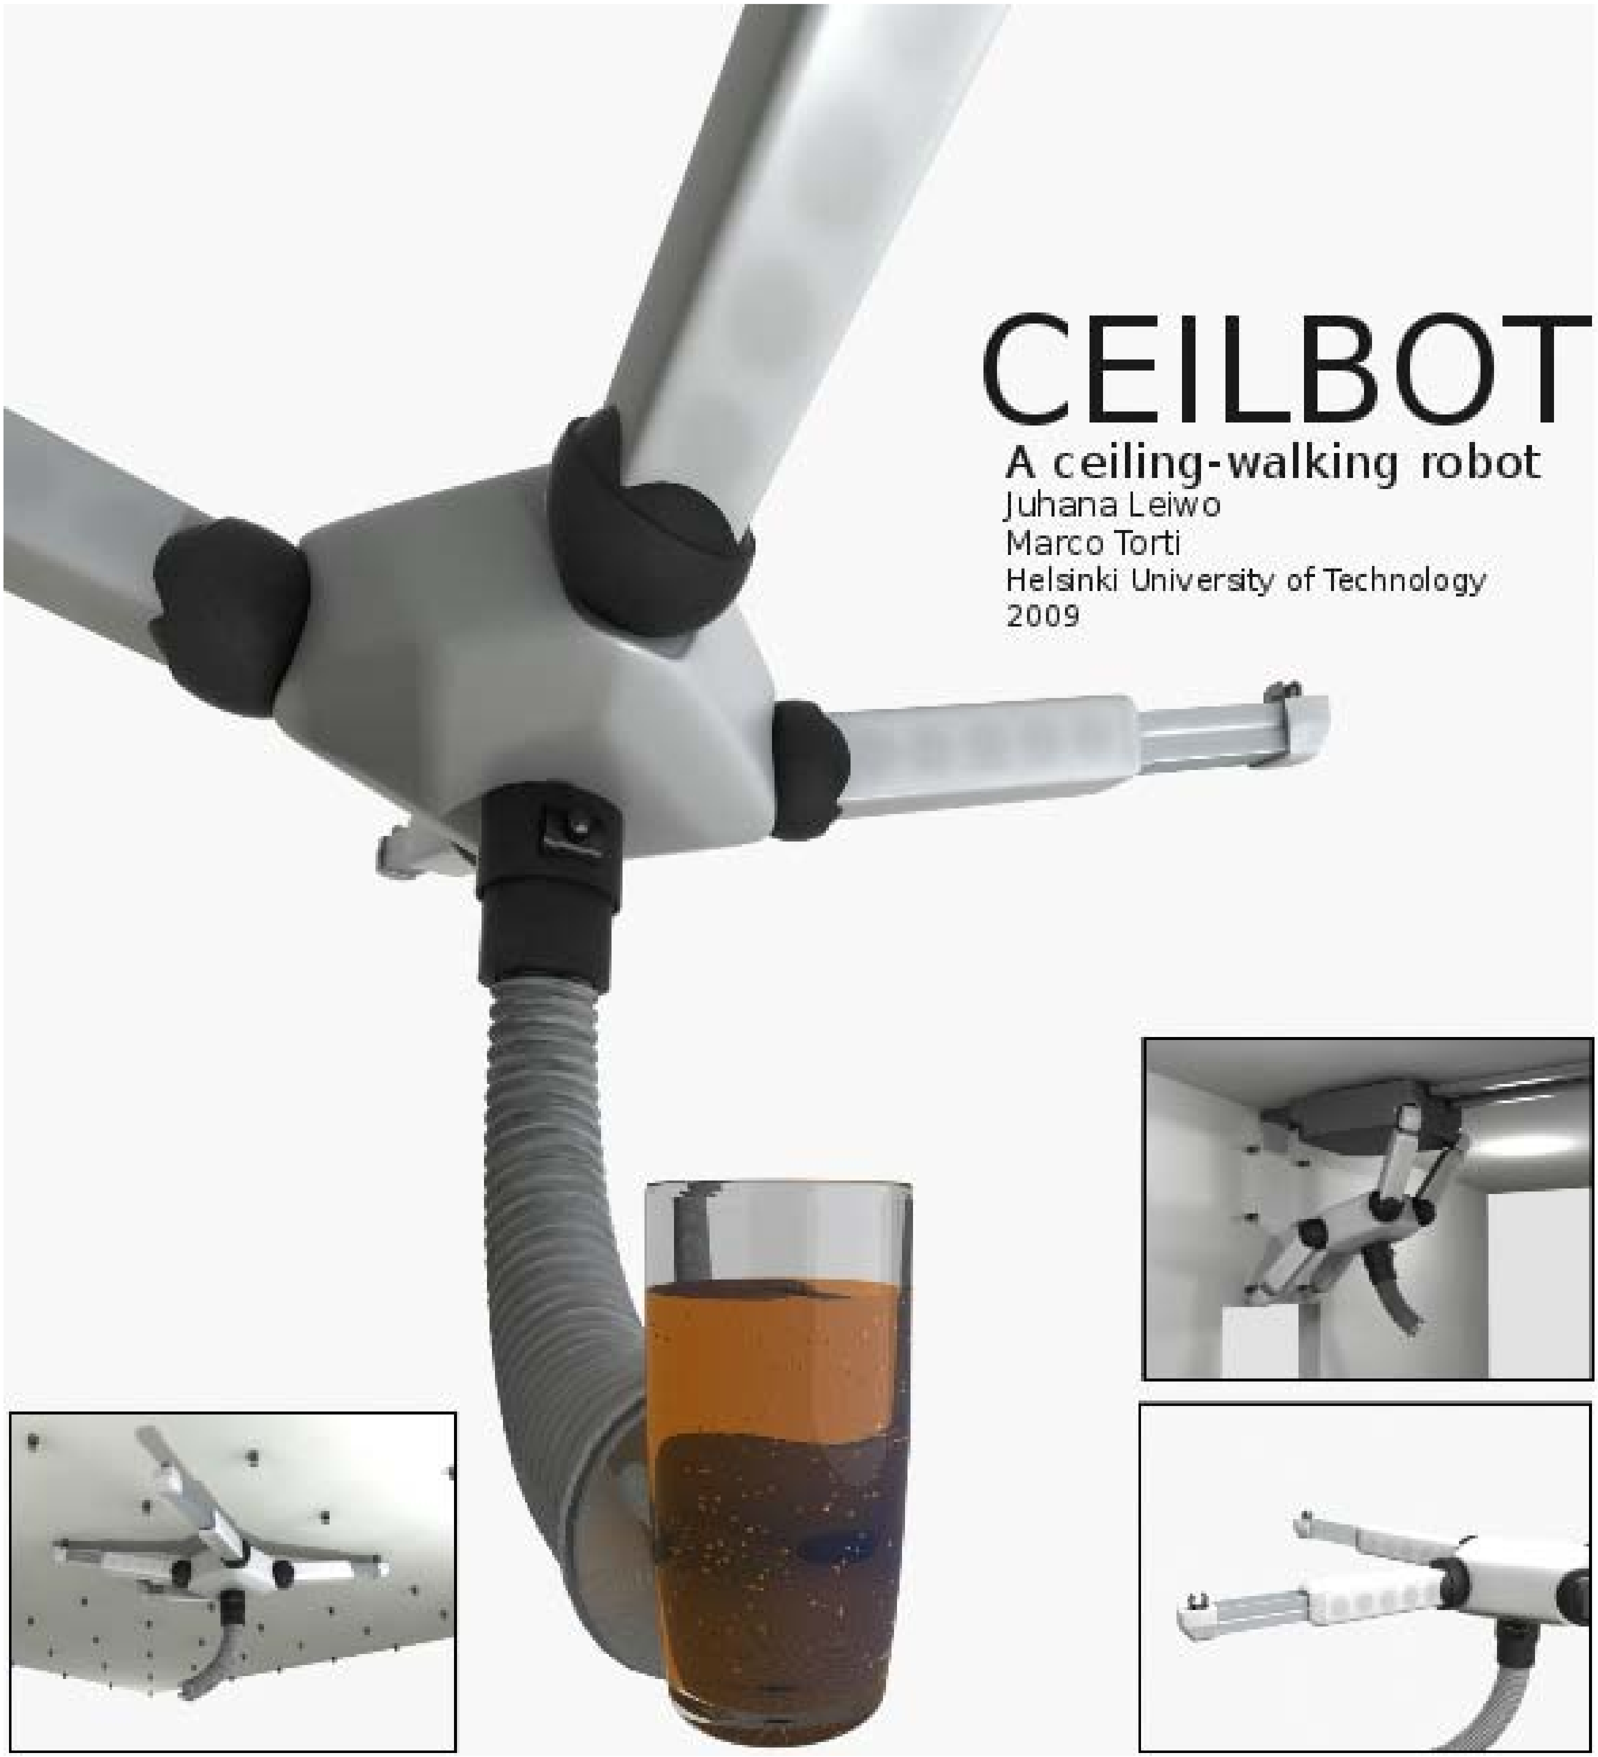

The project

The Ceilbot project is a study and research project organized at the Helsinki University of Technology. The aim of the project is to design and prototype a multifunctional robot which takes advantage of the free uncluttered space at the ceiling of a building's interior. The main motivation in designing a ceiling-going robot is to avoid the many obstacles at floor level in a human-inhabited space. The robot could therefore be built simpler and perform tasks which couldn't previously be robotized.
The first part of the project took place in the fall term of 2009. The participating students worked in groups of two and progress was reviewed each week in a meeting where each group held a presentation of their work.

General challenges and problems

The ceiling robot is intended to work in an environment designed for humans and it is supposed to be possible to install into a space not designed specially with the robot in mind. The challenges of operating must therefore be assessed in the robot's design alone, not by tailoring the operating environment to suit the robot. From a robot's point of view human living quarters are full of constantly changing elements that are difficult to categorize. Generally the only things that can be assumed to stay in the same place are the walls, floors, ceilings, doorways and windows.
Even though moving on the ceiling solves the problem of navigating among furniture and other objects lying around the floor, there are other issues to be solved. Perhaps the most obvious challenge is how to attach the robot on the ceiling securely so that it can still move around freely. Another point to consider is that even ceilings are not completely obstacle-free. There are many kinds of fixtures, air-conditioning pipes, lamps, wires, sprinklers and projectors that the robot should avoid collision with. Doorways also pose a design challenge, since a doorway is almost always lower than the ceiling and the ceiling can be on different levels on different sides of a door. Doors also need to be able to open and close.
Determining what the robot should be able to do and finding out how it might perform these tasks is a much bigger challenge than making it move around reliably. Many of the tasks discussed on the course are possible using currently existing technology. Others require a lot of research to become feasible. A good place to start could be to design a sensor system capable of providing the necessary information and feedback for a multi-purpose robot.

Our subjects

In the beginning of the course each team selected a specific application of the robot on which they concentrated. Instead of developing another form of the robot, we decided to work on some aspects of the design which would be common for all suggested applications. We chose to develop a location and mapping system and to study different methods of locomotion and attachment of the robot on the ceiling. Additionally we studied what different types of sensors the robot would need and what sort of safety protocol it should employ.

The attachment and locomotion system

The requirements for the attachment and locomotion system were that it should first of all keep the robot securely attached to the ceiling even in case of power failure, and allow it to easily move around and reach all parts of its operating environment. Secondly, the robot should be able to avoid obstacles like ceiling lamps and other fixtures, and preferably also pass through doorways. Initially several different types of system were discussed. Systems using magnets, suction cups or vacuum pumps were discarded as too unreliable, too weak and, except for suction cups, too energy-consuming. Rails were also left outside the scope of our group, because another group was already designing a railed system.

In the end, we decided to present two separate systems, one which uses chains and another which uses tactile limbs.

In the chain version, the robot is suspended by four chains, two on each side, from two actuators, one on each side of the ceiling of a rectangular room. The actuators are mounted on rails and are able to move along the length of the room. The robot is positioned laterally by moving the pair of actuators to cover one axis, and by adjusting the length of the chain pairs on each side to cover the other lateral axis. The vertical position is adjusted by giving slack or reeling in the chains on both sides. Two chains on each side allow the robot to resist torque perpendicular to the chains, and attaching the chains to moving actuators leaves more free space than hanging the robot from cables at each corner of the room. Chains are also stiffer than wires, which helps manage the pendulum motion of the suspended robot. Primarily the pendulum motion is to be cancelled by counter-moving the actuators with help of an acceleration feedback system. This system would be fast and relatively simple to control. It would allow the robot to pass under ceiling-mounted obstacles and reach all parts of a rectangular room easily provided there's no high furniture in the middle part of the room. It would also be easy to reach things that are low down without having to make the robot’s manipulator arm very long. This system would however be restricted to one room, because it can't pass through doorways. High pieces of furniture like bookcases and cupboards also limit the use of this system. (Or vice versa, this system limits the use of high furniture.)

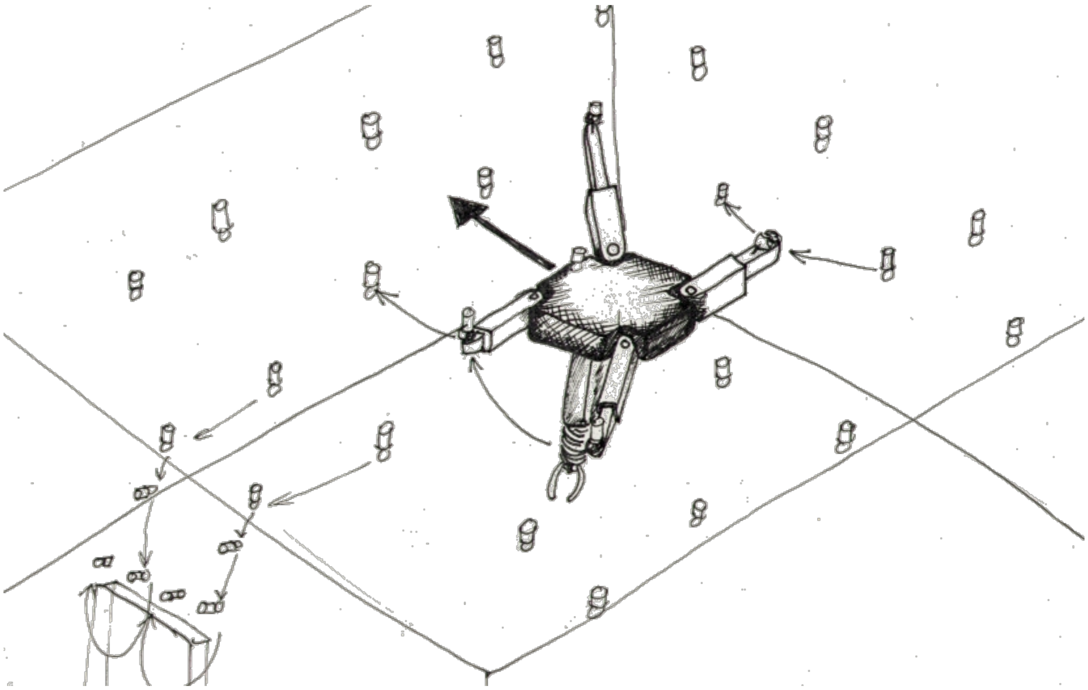

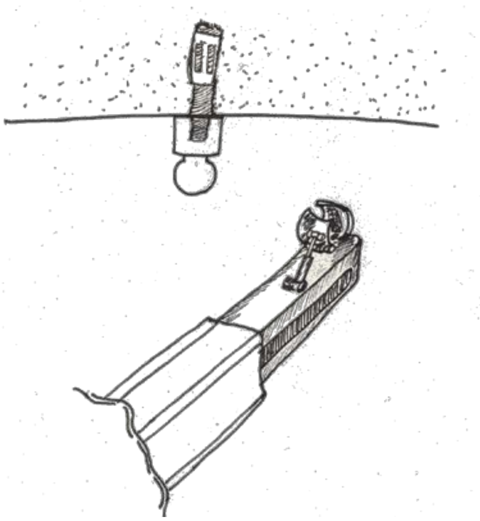

In the other one of our proposed systems, the robot uses four arms to grip metal anchor pins mounted on the ceiling. The arms are telescopic and they are attached to the robot with a two degree-of-freedom joint. The tip of each arm has an actuated gripping claw. The pin and claw are shaped so that when the claw is fastened to the pin, the pin is able to swivel like a ball joint. The pins are mounted on the ceiling in an evenly spaced grid. The robot moves around the room climbing from pin to pin with three of the arms fastened while one of them is changing grip.

The advantage of this system is that it can be mounted on almost any shape of ceiling. Passing through doors is also possible, if anchor pins are mounted on the wall above the door. Using this system, the robot would not be restricted to moving in the ceiling, it could use the walls as well if needed.
The downside of the climbing robot is complexity. The arms must move quickly and precisely for the robot to not be frustratingly slow. Also, it is probably not sufficient to rely on preprogrammed maps to locate the anchor pins, but some kind of active target-finding system must be used.

We also presented an idea of combining the two systems. The robot could use the climbing arms to attach itself to a mobile platform, which in turn uses the chain system to move in a room. This could be useful in an environment with one large room and many small ones. The large room would be equipped with the chain platform and the small rooms would only have anchor pins.

The mapping and locating system

Our other objective was to study the robot's mapping and location system. We presented an outline for a system in which the robot would use a layered data structure for mapping its environment and locating itself.

The idea of the layered map structure is to store things in the robot's environment in different data layers based on how often they change. The robot can be expected to be installed in a single operation environment for a long time. There can also be expected to be people in this operation environment, and people change things. Not everything changes, however. It can be safely assumed that the positions of walls, doorways, ceilings and such do not change at all, and much of the furniture remains in the same place for a long time. The mapping data structure would have one high-level map, which would contain the least-changing elements, one medium level map, which would contain elements that change seldom but aren't completely fixed, and a low-level map which only contains the most recently measured data. Also the frequency of change for different regions would be stored in a separate data layer. The robot would build these maps by itself, starting when it is first introduced to a new environment and gathering more data as it works.

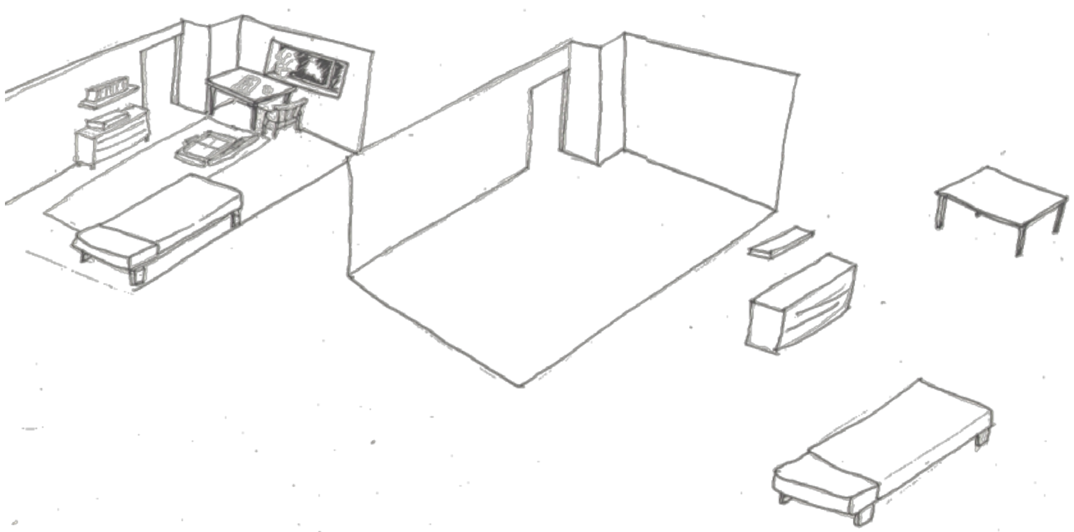

The high-level map

- The whole high-level map is used as a landmark when correcting location
- Points in the high-level map have a reliability factor which denotes how often the point has been averaged over.
- Further location-correction can use small spot scans of high-reliability areas instead of a full high resolution scan.
- The spot scans will also be averaged to the high-level map.
- To increase reliability evenly, spot scans will be taken from less reliable areas within reach in addition to the highest reliability portions.
- The averaging is done using a distance threshold between points in the old map and the new map.
- Also the unaveraged points are saved in the new map so that newly visited areas get saved. (With the exception of the medium-level points described later)þ
- The high-level map contains features that *have not changed.*

The medium-level map

- If a stored point is in a position which could be seen from both the old high-level map and the new scan, but exists in only one of them, it gets added to the medium level map.

- This means that only such points get added that were introduced or removed AFTER the previous scan of the area, that is, the robot can now see behind these points, or has previously been able to do so.

- Points added to the medium-level map get removed from the high-level map.

- This way the high level map will not contain features that have been added or removed during the robot's operation in the area.

- Features that already exist in the medium map will be re-added.

- The re-added features are averaged and given a reliability factor.

- If a feature that exists in the medium map is not found when later scanning the same area, it is removed from the map.

- Thus the medium-level map contains features that *have changed and may still exist.*

The change frequency map

- The change frequency map contains an intensity distribution that describes how often an area has experienced change.

- When a feature is added to or removed from the medium level map the intensity in the corresponding region of the change frequency map is incremented.

- The intensity values are made to decline over time so that areas that often change have higher intensities than areas of sporadic change.

With this sort of a mapping structure, the robot's navigation would become quicker as it grows used to its environment. It could plan its route ahead using the high-level map, and only perform quick spot-measurements to make sure the map still holds. A full-precision scan would only be needed at the moment of performing a task or in case the spot measurements don't have enough in common with the stored map.

The sensor system

For the sensor system, we concluded that the robot should rely on combining data from a number of basic senses instead of using proprietary sensors to support each individual function. We suggested the basic sensor system would contain proximity sensors, a multi-axis accelerometer and gyroscope, position information from the robot's servos, a video camera and a three-dimensional spatial scanner.

The camera and the 3d-scanner could be combined, if the scanning was performed using structured light triangulation. A planar-beam laser would be mounted at a fixed distance below or above the camera so that its pointing angle can be steered with a stepper motor. The laser beam's projection in the environment would be extracted from the camera image, and the information of the laser source's angle and position would be used to triangulate the projection image into three dimensions. The 3d data would be fed to the mapping algorithm described earlier.

For proximity sensing, we suggested ultrasound ranging, reflected light intensity measurement and touch whiskers. Ideally the robot would have all of its moving parts covered in a grid of proximity sensors. That way it could navigate in cramped spaces without knocking things over.

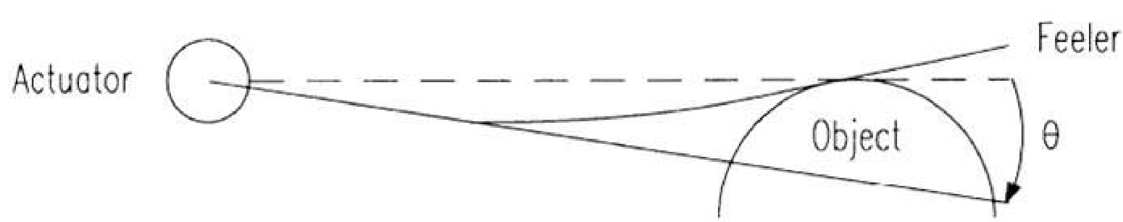

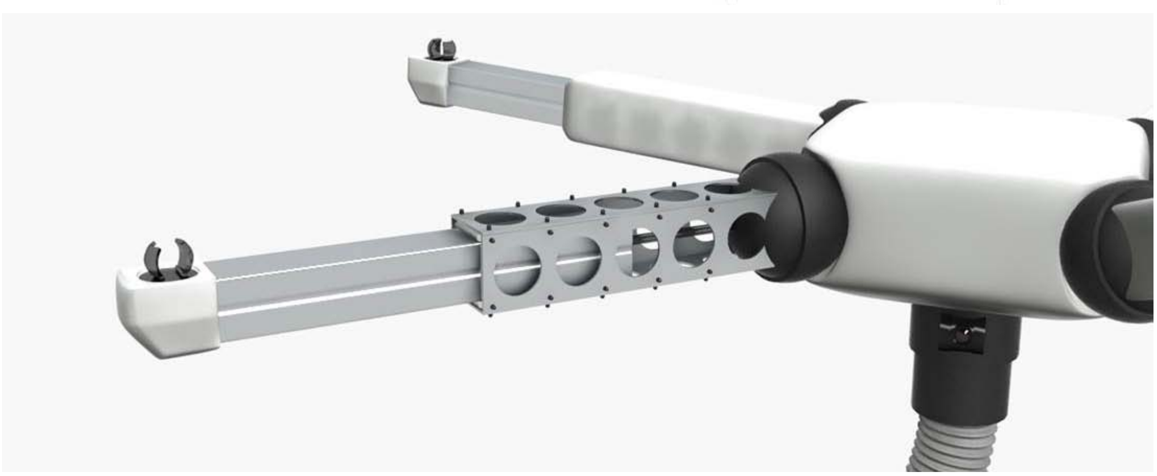

Because the robot is intended to operate in an environment with humans, it would help if it could distinguish between humans and objects. This would increase safety and allow more interaction with users. Primarily the robot would use facial detection from the camera image, but to aid this function we suggested using an infra-red heat sensor. The sensor would point in the same direction as the camera, and the robot would register the body heat of humans in its view. It could then try to stay out of the way or stop moving if people come close to it. The system would essentially emulate the function of a (very expensive) heat camera.

The robot could also be equipped with a stereo microphone, but more as an auxiliary sensor than part of the basic sensing system. The microphone would be used for detecting operators’ voice commands and, for example, signal beeps from other household machinery in case of tasks like emptying a dishwashing machine after it has finished washing.

Safety of operation

The basic rule is that the robot should not in any way disturb the normal activity of people in its operating environment. It should itself attempt to keep out of the way of people or, failing that, avoid hard impact in case of collision. The sensor system designed earlier already carries most of the sensors required for this.

In addition to the proximity sensor system on the robot's body, the manipulator arm could be equipped with extra touch sensors (whiskers). An unintended bump would then cause the manipulator to stop or even flinch away from the stimulus.

An obvious safety feature to have is an emergency stop switch. Hitting the switch would cause the robot to stop immediately. Besides a switch, a voice command could also be used for emergency stopping. Rather than a specific word, any loud sudden noise would be interpreted as a stop command. This way, even a barking dog could prevent the robot from doing something dangerous. (A thing worth considering is that most dogs are likely to bark at robots whatever it is the robots are doing. Perhaps the robot should have a constant false alarm rate algorithm to ignore unnecessary noise, or some form of human voice detection could be used.)

The robot must also avoid damaging property, including itself. Force feedback in the manipulator and locomotion arms would be used to apply the minimum amount of force required for an action. Picking up a pair of eyeglasses must be done with less strength than carrying a cast-iron frying pan, and the robot can't just be made so weak that it couldn't break the glasses anyway.

The robot should however not be strong enough to tear itself apart with its own actuators in any pose. Also the anchor pins on the ceiling must be able to support the combined weight of the robot and anything it is capable of lifting. At least two of the robot's locomotion arms must be attached on anchor pins at any moment (three arms are required for stability anyway) and the arms must maintain grip even if power is removed.
The robot is connected to electrical power through plugs in the gripping claws. Each anchor pin has a socket with both electrical poles and each claw has a plug. This way the power source is always connected when the robot is on the ceiling. The power supply to the anchor pins' sockets is low voltage DC, so that the robot doesn't need to carry a transformer. To ensure operation even in case of a power cutoff the robot carries a reserve battery with enough capacity for about five minutes of normal operation and one hour of operation for the control computer. Of course the “pin grid” could also be fitted with an Uninterruptible Power Supply (UPS) system, but an onboard battery works even if the pin grid itself fails.