

Current Research Projects of the Artificial Intelligence Group (2009)

Prof. Dr. Raúl Rojas

Robotic Projects

R1: Autonomous Cars	2
R2: Humanoid Robot Design	4
R3: Innovation Laboratory AutoNOMOS.....	5

Information Appliances and Image Processing

IA1: An Information Appliance for the Blind.....	6
--	---

Biology Inspired Systems and Bionics

B1: Reading Glasses for the Blind	7
B2: Building a Robotic Bee	9

Ambient Intelligence

A1: Intelligent House Project	10
-------------------------------------	----

Human-Computer Interaction and E-Learning

E1: Classroom of the Future.....	12
E2: Pen-based authoring tool for E-Learning	13

History of Computing

H1: Digitizing Konrad Zuse's "Nachlass"	14
H2: Reconstruction of the addition unit of the Z3	15
H3: Reconstruction of the Z1	16

Startups	17
-----------------------	----

Old Research Projects	18
------------------------------------	----

Robotic Projects

R1: Autonomous Cars (2006-2009)

<http://www.spiritofberlin.eu>

We have been building autonomous robots for many years. In the last few years, we started developing *full-scale robotic vehicles*. Our group took part in the DARPA Urban Challenge competition, held in November 2007 in the USA. We reached the semifinals of the contest. Our vehicle has been operational since January of 2007. We are now driving the vehicle in our new testbed: Tempelhof Airport in Berlin.



Our robotic vehicle at the Urban Challenge competition in Victorville, CA (2007)

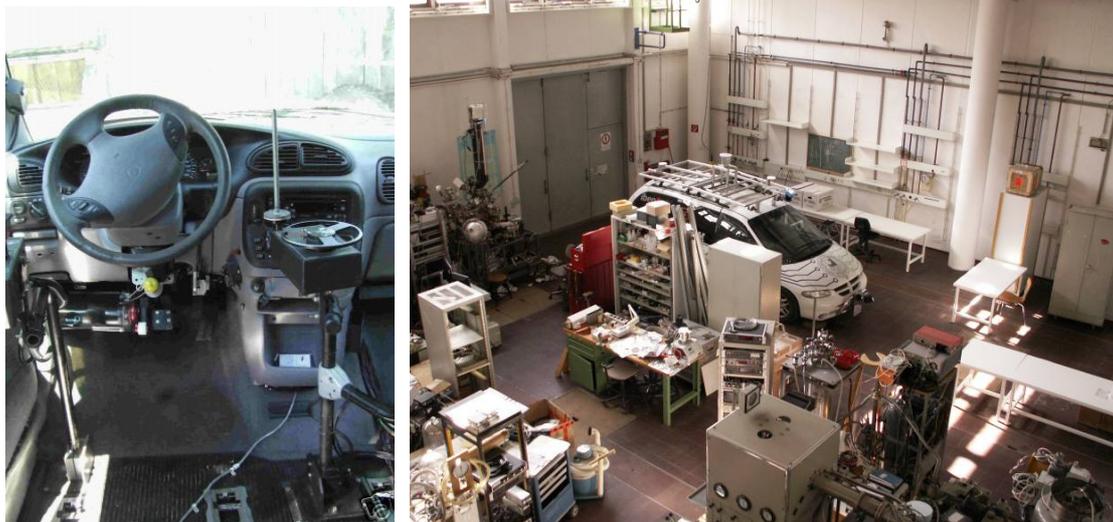


Spirit of Berlin driving in Tempelhof Airport (2009)

Our vehicle has the following sensors:

- a) Frontal laser scanner for detecting vehicles on the road,
- b) Velodyne 3D-Scanner on top of the vehicle,
- c) Video cameras on top of the vehicle,
- d) Ultrasound sensors in the bumpers,
- e) GPS navigation system with submeter accuracy.

The car is controlled by two IBM servers with four cores on each board.



A view of the cockpit of our vehicle, adapted for drive-by-wire, and of our laboratory at FU Berlin

Future components and sensors for our autonomous car will be cheap, modular, and reusable on different kinds of robots.

We are working on several subprojects:

- a) Computer vision for the road. The computer vision should be able to recognize vehicles and the lane lines. We have a prototype for a car recognizer developed with machine learning techniques. In a cooperation with the company Hella-Agliaia we are now using their computer vision development software.
- b) New method for tracking objects detected with the laser scanner. At the moment of this writing, we detect objects with the laser scanner in every frame but we do not track them more than a few seconds. It is not trivial to track scanned objects (due to occlusions).
- c) Combination of the computer vision with the data from the laser scanners, in order to implement a “virtual Lidar sensor” with higher resolution.
- e) Development of a new kind of 3D-Scanner. We used a Velodyne 3D-Scanner during the Urban Challenge competition. This is a very expensive, yet powerful unit. We are developing our own 3D-Scanner using a special mirror in collaboration with MicroEpsilon, a LIDAR company.
- f) We are investigating also high-level control strategies for autonomous vehicles based on reinforcement learning.

In the future we would like to have one single 3D-laser scanner and videocameras as the main sensors. We would also like to be able to drive without an expensive GPS system. The navigation system should be like the ones used today in consumer cars. We are working on such a project in cooperation with Hella-Agliaia.

This project is being pursued by five doctoral students: Fabian Wiesel, Tinosch Ganjineh, Mohammed Fraheer, Miao Wang, and David Latotzky with funding from the Ministry of Education and Research (BMBF). Additional funding has been requested for instrumenting a second autonomous car in 2010. Five undergraduates participate also as members of the development team. In a related project, my PhD student Colin Bauer from Carmeq AG, is developing cognitive models of human drivers. We will test these models using our robotic platform.

R2: Humanoid robot design (2006-2009)

<http://fumanoid.mi.fu-berlin.de>

In this project we have developed three generations of biped robots able to walk and balance on two legs. The stereoscopic head developed in another project has been mounted on this robot.

The main problem when designing bipedal robots is the large amount of degrees of freedom available once several motors are used to simulate the human joints. In our project we developed special PIC circuits of small size that help us to keep the complexity of the task in check. The controllers used for the legs are fuzzy controllers based on a set of fuzzy rules. The robot is modular.

We have investigated gaits based on keeping the center of mass under the contact points of the robot. One way of achieving this is to use a pendulum to simulate the movement of the upper body. The gaits we obtained allow the robot to move forward, backwards, and turn around.

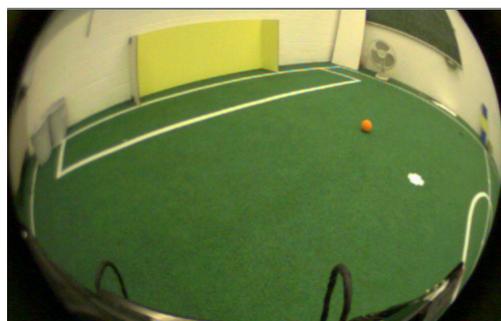
In the next stage of our project we are now improving the speed of the robot and are investigating different types of gaits, including the passive dynamic walking approach. It consists in moving forward the center of mass, accelerating the robot, and avoiding a fall by synchronizing the leg reflexes with the balancing movement.

We are developing a fully self-calibrated stereoscopic computer vision for the robots.

Our humanoid robots won third place at the RoboCup competition held in Atlanta in 2007, they won the Iranian Open in 2008, and obtained recently third place at the German Open 2009.



Our humanoids in RoboCup 2007 (Atlanta) and in Hannover, German Open 2009.

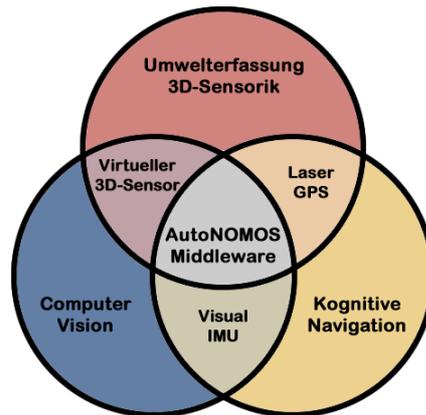


The world as seen by one the robot's cameras

This project is being coordinated by Hamid Moballegh, and Gretta Hohl with funding from Deutscher Akademischer Austauschdienst (DAAD). Ten undergraduate students participate in the team.

R3: Innovation Laboratory AutoNOMOS

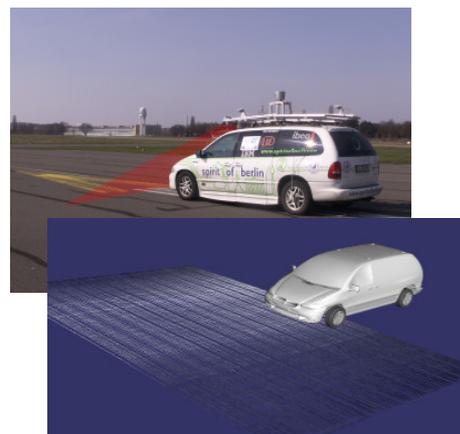
The first phase of this project is being funded by the Ministry of Education and Research (BMBF). The project consists in developing the middleware for the connection of different kinds of driver assistance systems. Until now, such systems have been sold as black boxes. They consume information from sensors delivered through the CAN bus, and provide a result, usually just a warning signal, but do not communicate with each other. In our project we are developing a kind of operating system for three kinds of software modules: computer vision software, Lidar and Radar sensors, and for low-level and high-level control.



The project has been divided into three main software groups (CV, 3D sensors, Navigation) and three interconnected modules: visual IMU, Laser GPS, and virtual 3D sensors. The core system will implement new protocols just adopted for automotive engineering. The core system will be based on the concept of virtual sensors, which will allow future driver assistance systems to be mounted on cars in “plug & play” manner.

The complete system will be first used to instrument special vehicles. In one of our applications, an autonomous vehicle will measure and check the state of the runway at an airport periodically (this is currently done by persons performing a visual check). In another application the Tempelhof Airport will be circled by an autonomous car instrumented for security checks and surveillance. Other applications are planned for the next two years.

Dr. Andreas Stopp (previously at Daimler) is directing this project. Sebastian Gaissert, Martin Storre, and Tinosch Ganjineh complete the core team. Further personal will be hired in the second phase of the project, to be started in 2010.



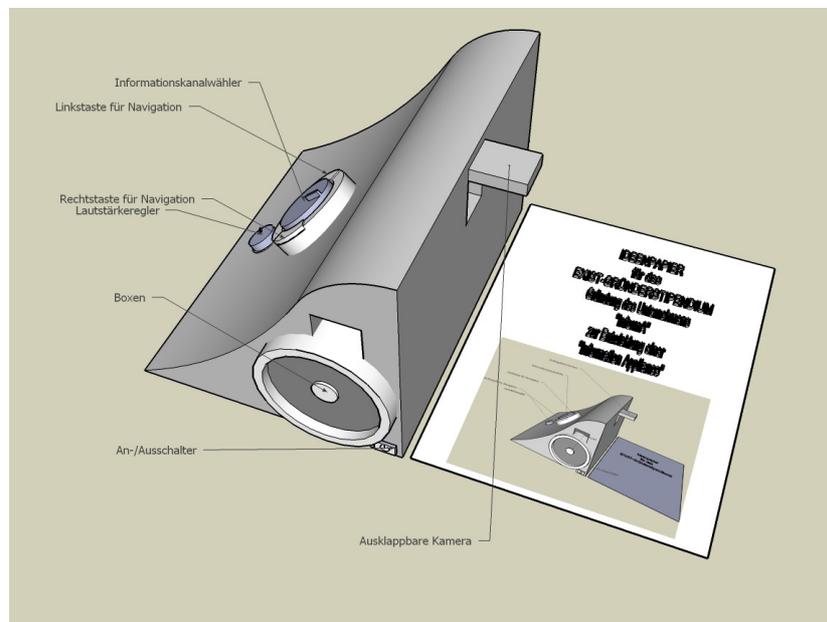
Inspecting the runway at Tempelhof Airport

Information Appliances and Image Processing

IA1: An Information Appliance for the Blind (2008-2009)

In the years 2007-2008 we developed a reading device for blind persons. A digital camera captures an image of a document, the computer recognizes the layout and the text in the document, and reads it aloud. The device is portable, it consist of an embedded PC with a battery and a 5 Megapixel camera. The reading device is now being sold by the company Beyo GmbH in Potsdam. I received the Technology Transfer Prize “WissenWertes 2009” for the development of the machine.

We are now developing the next generation of the reading device which has evolved into a complete information appliance for blind and elderly persons. The picture below shows a diagram of the machine. As in the first reading device, a digital camera is used to capture documents, but now the recognition and OCR are done accessing a server through the Internet. The computer in the device is very small and uses low-power chips. Speakers in the machine are used to play the synthetic voice files sent by the server.



The main innovation is that a blind or elderly person can navigate the Internet using a rotating dial. The use of the device is similar to tuning radio stations using a portable radio. The stations, in our device, are services which the user can access. One channel or service is OCR for reading documents aloud. Another service is a call center which can help to configure the device, or reading fonts which are too small or difficult to recognize for the computer. Another information channel is the daily newspaper, which can be read by the computer. The user navigates through the different sections (organized as a tree) using the dial and pushing the knob for selecting a branch in the tree (in a similar way as with the old iPod). Additional services are: providing the time of the day, the weather prediction, hearing e-mail, dictating e-mail, tuning Internet radio stations, etc. The number of different services grows continually and does not modify the information appliance itself, only its usefulness.

Two PhD students (Roman Guilboud and Bastian Hecht), as well as three students are collaborating in this project. We have a collaboration with Pablo Vidales at T-Labs. This research center and ourselves will conduct a field trial in Berlin with 50 devices and the necessary DSL connections during 2009. The project is being funded by Deutsche Telekom AG.

Biology Inspired Systems and Bionics

B1: Reading Glasses for the Blind

Imagine having reading glasses with a small video camera attached. The signal from the camera is sent to a small sub-notebook computer. The camera is able to rotate because it is attached to a spherical joint, so that it can capture images as if it had a wide-angle lens. Imagine a blind person using these glasses. He or she sits at home, holds a newspaper in front of his or her face, and using spoken commands, orders the computer to read the text aloud. The computer first assembles a single picture from the several images provided by the camera and detects the boundaries of the book. It then recognizes the headlines and reads them aloud using a software speech synthesizer. The blind person then orders the computer to read the first article, the third, and so on.

Imagine the same blind person being able to walk around a building and order the computer to read the nameplates on doors. The camera actively looks for doors, captures the image, and proceeds to recognize the text written on them.

Imagine a blind person entering a restaurant and being able to read the menu, or the label on a bottle. Imagine a blind person able to ask the computer, with a spoken command, to look for windows or doors in a room. Imagine that the computer tells the blind person that it can detect five faces in a room, so that he or she knows how many people are there.

This is exactly what we want to achieve with our project. We want to build a portable reading system for the blind, which can also help in simple pattern recognition situations.

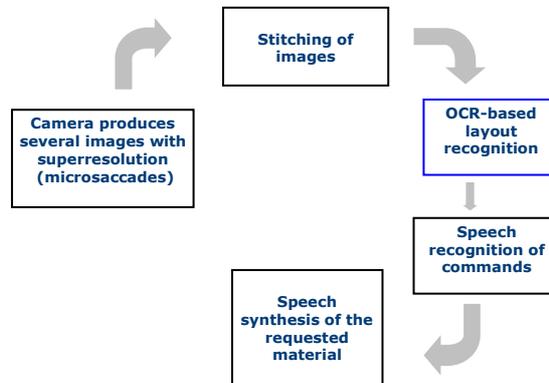
During 2005-2007 we conducted a project in which we developed the first version of Saccadic(i), our vision of intelligent reading glasses for the blind. We did the following:

- We applied piezoelectric elements to move the imaging chip of a video camera.
- We developed super resolution software based on the drizzling method, which improves the image resolution fourfold.
- We manually controlled the camera to take several images of a book or newspaper.
- We integrated a commercial stitcher in our system, to provide the images for the OCR.
- We wrote layout recognition software of our own using the OCR Fine Reader recognition engine library.
- We wrote a navigation interface for spoken commands.
- We integrated a commercial SpeechXML-based speech recognizer with the system.
- We packaged the complete software into a JVC laptop with 1 GHz and 1 Gigabyte of memory.

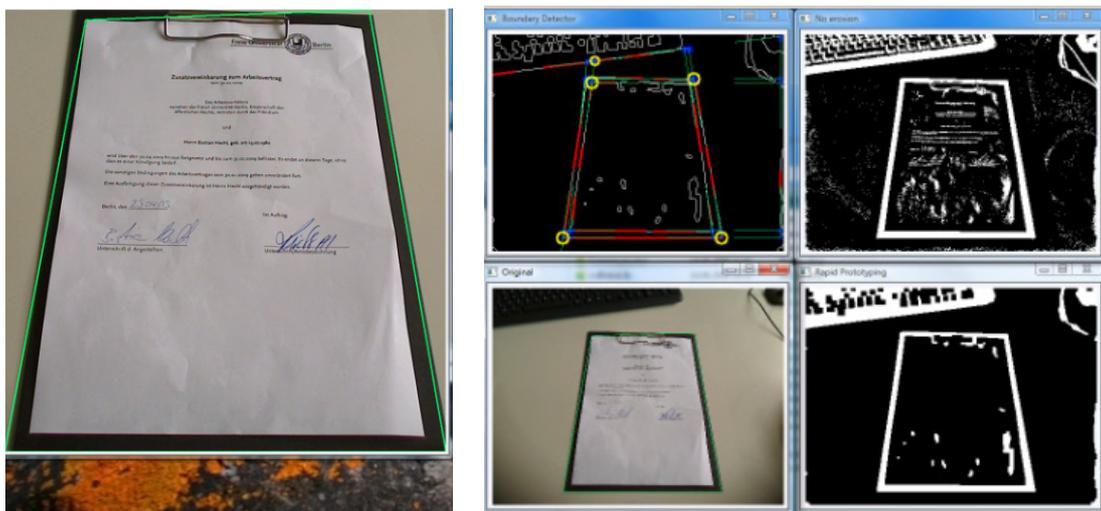
The system works and the first public demos have been made. The major shortcoming of the current system is that it is slow, stitching requiring several seconds, and it is not portable. The camera system is too heavy and the laptop is still too big. In our project we have to:

- Miniaturize the system.
- Speed it up to real time.
- Make it robust against changes in lighting.
- Improve the image recognition accuracy.
- Integrate more features.

The diagram shows the steps followed in our current Saccadic(i) prototype. The video camera is used to take images and increase their resolution using drizzling. Several images of a book or newspaper, displaced, are taken. We use manual control. The stitching software automatically produces a single image, which is passed on to the layout recognition software. This is based in the Omnipage library, for which we have a license. The laptop accepts spoken commands and can read the text using speech synthesis.

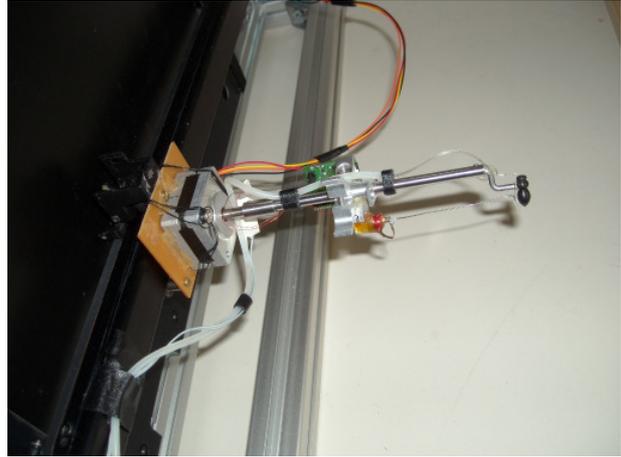
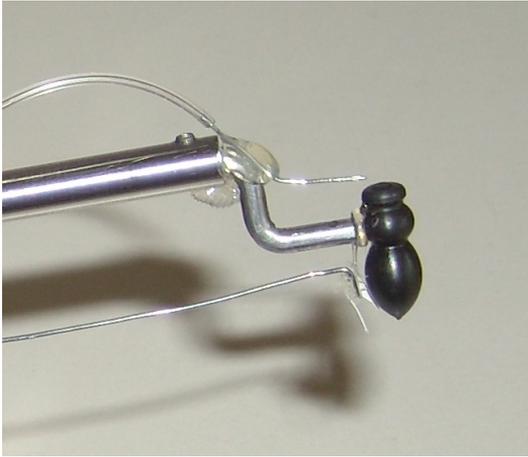


Dataflow in our prototype system



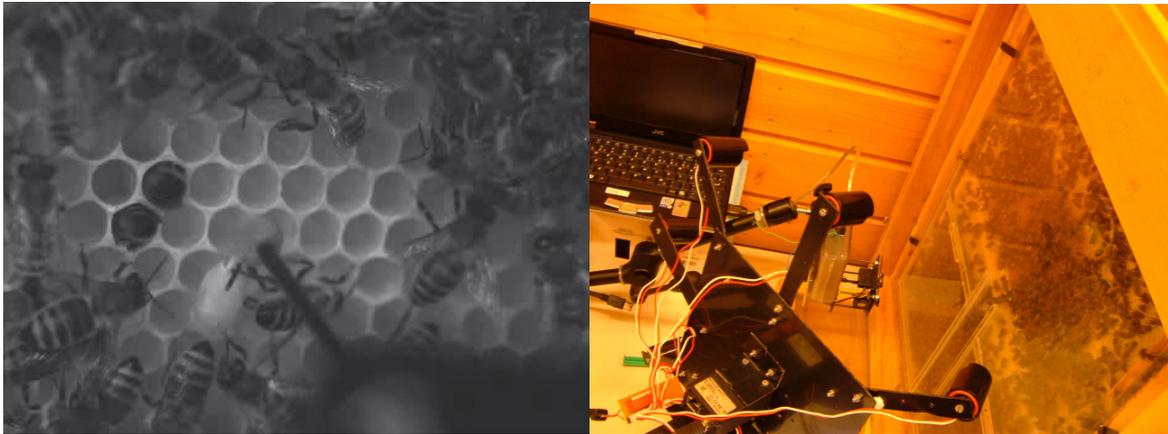
A document captured by the portable camera. Several processing steps can be seen on the right.

This project was a finalist in the 2006 “Bionics Ideas Contest” conducted by the Ministry of Education and Research. The new phase of the projet is being funded by the DFG. Three doctoral students (Bastian Hecht, Roman Guilboud, and Marte Ramirez) and one postdoc (Marco Block) are working on this project. Several students are writing diploma thesis dealing with different aspects of the system.



Two views of the robotic bee assembly. The wings are on place and a tube for providing sugar to other bees.

The first mechanical version of the robotic bee has been tested in the beehive. The photographs below show the experimental setup (to the right) and the robotic bee in the beehive.



Robotic bee in the beehive (left) and experimental setup (right)

The figure below shows an antenna attached to a real bee. The antenna reflects 9 GHz pulses at a harmonic frequency of 18 GHz, which can be tracked by a radar unit. In the experiment, we will let the robotic bee dance on the beehive. Bees leaving the beehive will be marked with an antenna, and their path will be followed with the harmonic radar.



A bee with a harmonic radar antenna

Hamid Moballegh (PhD candidate), Tim Landgraf (PhD candidate), and around seven students collaborate in this project with the Biology department at FU Berlin (Prof. R. Menzel). The first field tests of the robot are programmed for the summer of 2009. The project is being funded by the DFG. New funding from the Bernstein Center for Neurobiology has been granted for 2012-2014, for the development of insect-like robots.

Ambient Intelligence

A1: The Intelligent House Project

We are building a house in Berlin featuring the currently most advanced house automation systems. All electric appliances in the house, all lights, and the window shutters can be controlled through the Internet and using an iPhone. Several robots perform different duties in the house, from cleaning to providing security services.



These are some of the house highlights:

- Automatic control of all switches and lights in the house through the Internet or iPhone
- Remote energy consumption metering in real time
- New energy sources for heating: solar power and heat pump
- Electronic office with videoconference equipment for sending lectures to the university
- Electronic kitchen with all devices connected to a wireless control panel
- Wireless alarm system with motion, smoke, vibration detectors and remote communication
- A mowing robot takes care of the garden
- A guard robot patrols the house
- Two cleaning robots for the floor
- A computer and storage server for the house
- Two tables with embedded large LCD screens for games or reading the newspaper
- Virtual fitness room

In the garage the car will exhibit an electric car, a hydrogen fuel cell car, and a solar powered bicycle, among others.

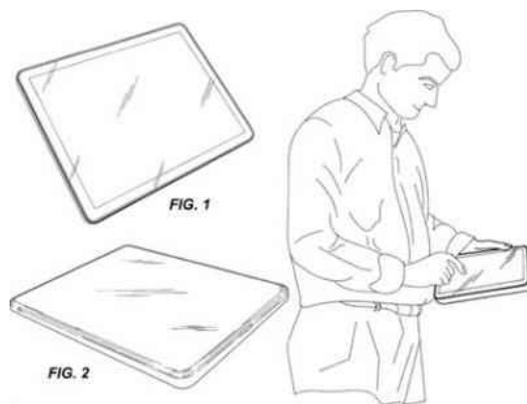
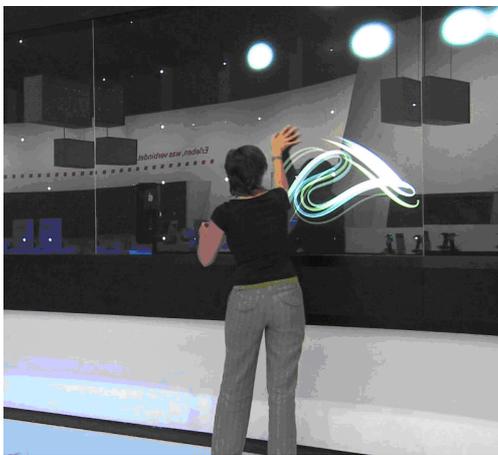


Several students are collaborating in this project writing the software for control, and for the robots. David Seelbinder is coordinating the project. Martin Klein is writing the media server. The house is being built with private funding and company sponsoring.

Human-Computer Interaction and E-Learning

E1: Classroom of the Future

In this project, a continuation and upgrade of our old E-Chalk system, we are investigating the hardware and software needed for the classroom of the future. Electronic media will play an important role in such a classroom. We want to substitute the traditional chalkboard with a large interactive multitouch screen. The image below shows the kind of large displays that we are aiming to (the installation was done by one of our students in the project for the International Radio Exposition in Berlin in 2008). We want the students to be partners of this collaborative educational effort. Therefore, they will be provided with interactive pads. Students can post material to the larger screen, can capture their own annotations and use them as a second layer of meaning superposed on the lecture materials. They can share and co-develop annotations.



The screen will be intelligent: handwriting will be interpreted and gestures will launch or stop specific programs or intelligent agents. The computer screen will thus transform into a living interaction medium.

We are also experimenting with 360-degree videos of the lectures. This video material will be superposed on the digital media. Since we capture the complete lecture room, omnidirectionally, we are investigating the social editing (that is the collaborative editing) of the videos. The image below shows some of the videos we have captured with a six cameras omnidirectional video system and the partial views which can be generated by the software.

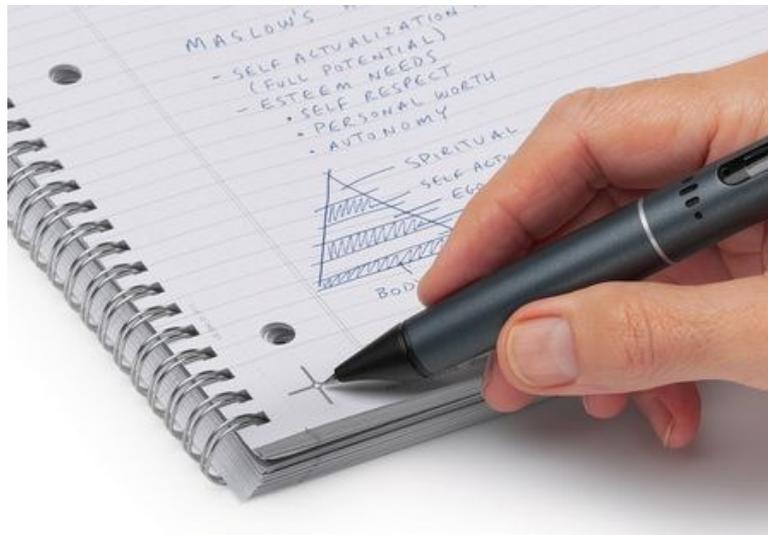


Several persons are collaborating in this project: Dr. Ernesto Tapia (postdoc) and three more students. The project is funded by FU Berlin.

E2: Pen-based authoring tool for E-Learning

PowerPoint slides are very popular among lecturers because they minimize errors during class, let the lecturer keep his or her discourse coherent, and because it is very easy to prepare the slides. We have been investigating for many years the opposite scenario, when the lecturer writes directly on the blackboard delivering a science class as a kind of performance. The blackboard is still used by mathematicians, physicist and natural scientists in general, because it allows thoughts to be developed at a pace which allows the students to grasp the material and ask questions.

In this project we are developing an authoring tool for the electronic blackboard. A special digitizing pen is used to capture pieces of the lecture (such as definitions, notation, diagrams, or even proofs), storing them as sequences of strokes which can later be used for a class. The user just writes on plain paper his or her notes, before the lecture is given, and the complete set of strokes can be pasted on the electronic blackboard at a later point, on demand.



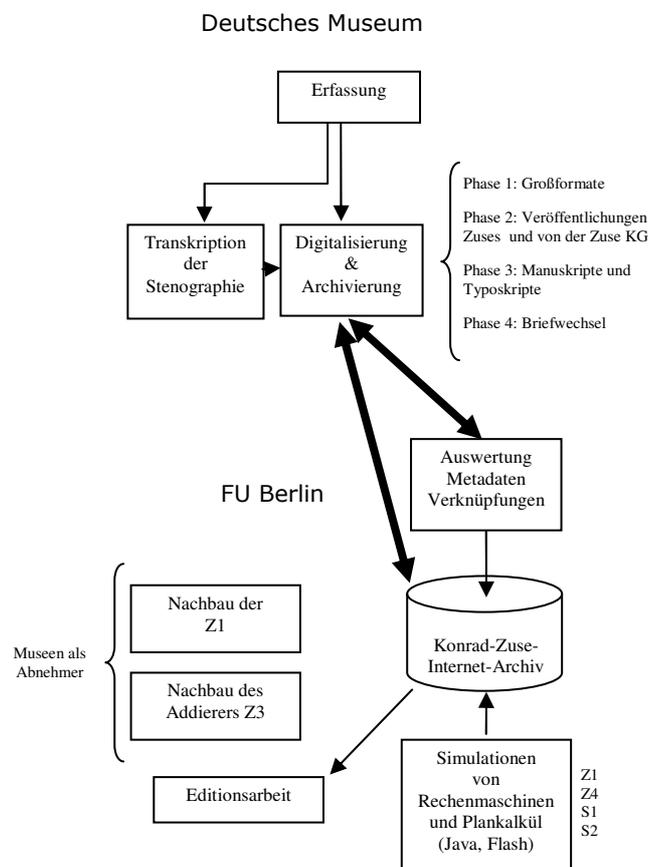
Dan-El Neil Vila Rosado (PhD candidate) is the leader of the project. FU Berlin is financing the project with an innovation fund scholarship.

History of Computing

H1: Digitizing Konrad Zuse's "Nachlass"

The Klaus Tschira Foundation bought Konrad Zuse's papers ("Nachlass") in 2005 and donated the complete collection to Deutsches Museum in Munich. The acquisition contract stipulates that I should be given unhindered access to the archive with the purpose of doing the scientific evaluation of Zuse's documents.

We have been working with Deutsches Museum towards digitizing the collection of documents. The image below shows the work-flow of the project. Deutsches Museum is doing the digitizing and transcription work, while FU Berlin is providing the expertise for the classification and possible edition of the documents. The digitized documents will be put on two servers, one in Munich, the second in Berlin.



Part of the work consists in writing Java simulations of the first Zuse machines: the Z1, Z2, Z3 and Z4.

This work is being done by PhD candidate Till Zoppke (DFG funding) and our collaboration partners at Deutsches Museum.

H2: Reconstruction of the addition unit of the Z3 (2009)

In 2010 we will be celebrating the 100th anniversary of the birth of Konrad Zuse. In 2001 we reconstructed the complete Z3, the computer built by Zuse in 1941 (see image below). For the 100th anniversary we just completed a “desktop” reconstruction of the addition unit of the Z3. The board works with 5V and uses small relays. The circuit is the original one described in the patent application of 1941.



The Z3 rebuilt in 2001 for the sixtieth anniversary of the patent application

The addition unit of the Z3 contains two registers, uses two's complement arithmetic, can add and subtract. The original Z3 worked with 22 bits and two addition units: one for exponents, another for mantissas of a floating point representation.



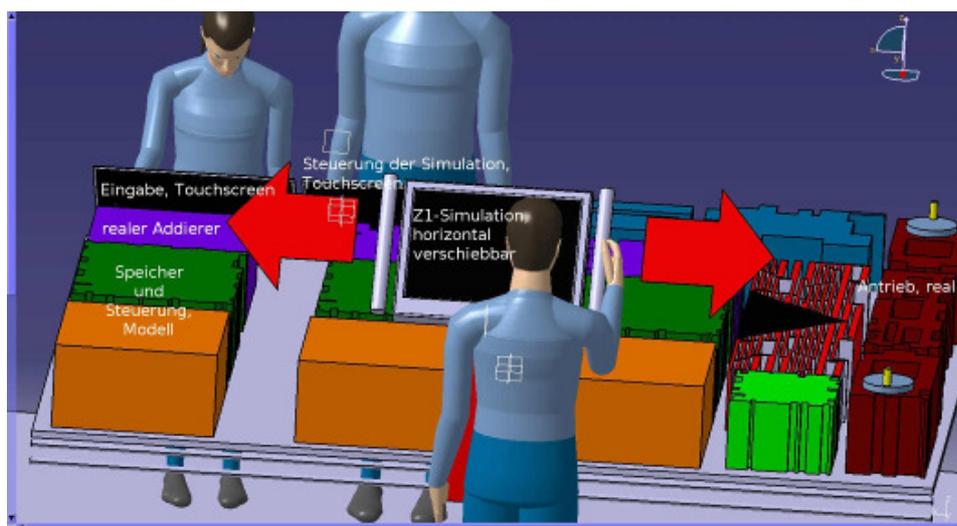
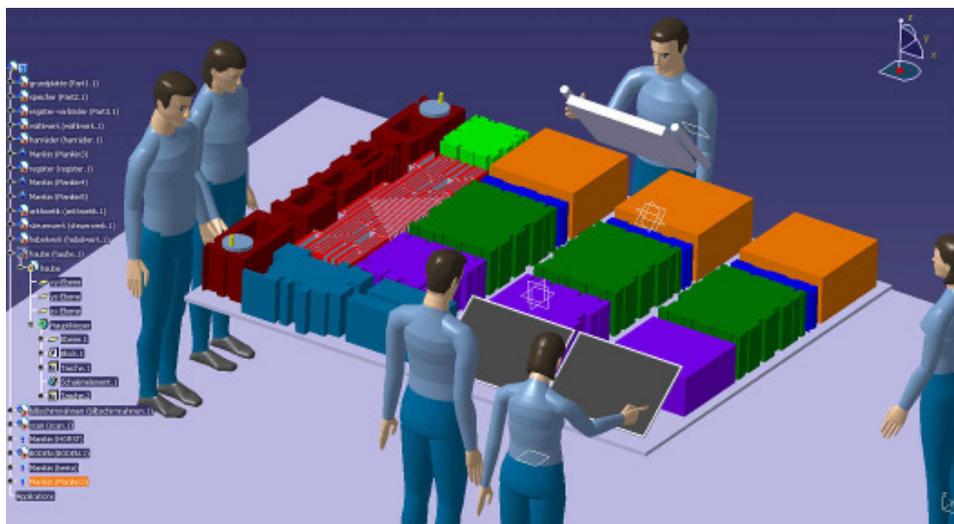
The new built addition unit of the Z3

The project has been funded by the Klaus Tschira Foundation. The design was made by Frank Darius, Georg Heyne, Alexander Warth and Raul Rojas.

H3: Reconstruction of the Z1

After having rebuilt Konrad Zuse's Z3 computer in 2001, we have been studying the circuits and mechanical assemblies of the Z1, Zuse's first computer completed in 1938. The Z1 is the predecessor of the Z3, has a much more complex structure, and it is safe to say that is not currently understood by anybody in the world. Accordingly, the first part of our project consists in doing a thorough study and simulation of all circuits in the Z1.

The second part of the project consists in rebuilding a subset of the mechanical assemblies and complementing the result with an "extended reality" display. The visitors of the museum where the reconstruction is to be housed will be able to see some assemblies moving and computing results. The missing assemblies will be simulated in 3D and their operation will be visible in a computer display. The users will be able to "see through" the mechanics in order to understand the way the Z1 worked. The figure below shows a CAD image of the intended result. The digital displays will be semitransparent and will be mounted close to the mechanical assemblies.



This project is being conducted by Till Zoppke (PhD candidate), and is being financed by the Klaus-Tschira-Foundation, and the Heinz-Nixdorf Foundation. The machine will be installed, when finished, at the Heinz-Nixdorf MuseumsForum in Paderborn.

Startups

I have been involved in various startups derived from our work at the university.

Beyo GmbH

I was a co-founder of Beyo GmbH, a company specializing in reading devices for the blind, together with Oliver Tenchio and Cüneyt Göktekin. The first product has been on the market since 2008. I received the Technology Transfer Award from the Technology Foundation Berlin in 2008 for this development. The company has gone through two rounds of venture capital funding since 2007. I sold my shares to investors in 2008.

Here I am receiving the Technology Transfer Prize (and check)



Appirion UG

I am a co-founder of the company Appirion (2009) together with Miao Wang and Tinosch Ganjineh. Appirion (Appsolute Touch before) is now selling two applications for the iPhone, the first is a translation tool, the second a lifestyle/game application.



WohnoSapiens UG

I am a co-founder (2008) of WohnoSapiens a company specializing in intelligent house automation. The company will start offering software tools in 2010.

OMQ GmbH

I am the entrepreneurship coach for the new company OMQ in Berlin. OMQ is developing an intelligent tool for remote servicing and diagnostic of software applications (OMQ Diagnostic). My task is to oversee the development of the machine learning tools.

Old Research Projects of the Artificial Intelligence Group

Robotic Projects

R1: FU-Fighters: Autonomous Soccer Robots

R2: Adaptive Vision for Mobile Robots

R3: Neural Coach for Autonomous Robots

R4: Stereoscopic Vision for a Biped Robot

Human-Computer Interaction and E-Learning

E1: E-Chalk: An Intelligent Chalkboard

E2: Software Architecture for a Smart Teaching Board

E3: Intelligent Audio and Video Processing for E-Chalk

E4: Recognition of Handwritten Mathematical Formulas

E5: Algorithmic Animation for E-Chalk

E6: Hardware Design for a Multimedia Classroom

History of Computing

H1: Konrad Zuse Internet Archive

H2: Reconstruction of Konrad Zuse's Z3 Computer

H3: Simulation and Reconstruction of the ENIAC

Old Research Projects of the Artificial Intelligence Group

R1: FU-Fighters – Autonomous Soccer Robots (1998-2006)

www.fu-fighters.de

We have been building autonomous soccer robots at the FU Berlin since 1998 and took part in our first RoboCup tournament in 1999. We were active in two leagues until 2006.

Our team, the FU-Fighters, was composed of two types of robots: small-size (18 cm in diameter) and medium-size (50 cm in diameter). The small robots are controlled by an off-the-field computer, which receives video images from two cameras positioned above each half of the 4 by 5.5 meters field. The computer sends the robots commands using a wireless link. Our medium-size robots are fully autonomous: they carry their own camera and a laptop, which processes the video images at 30 frames per second. Our medium-size robots use a special mirror that provides a 360 degree view of the field (omnidirectional vision). Both types of robots, small and medium-size, have omnidirectional wheels.

Our soccer robots are an ideal platform for investigating Artificial Intelligence issues such as behavior control and learning. We have experimented with adaptive algorithms for computer vision and with reinforcement learning for strategy. Our robots can learn from their experience during actual games.

The small-size FU-Fighters robots are the World Champions 2004 and 2005. They won second place at the RoboCup tournament three times, third place once, and fourth place once. They are four-time winners of the international German Open tournament, and European Champions 2000. The medium-sized robots won fourth place at the German Open 2004, and second place at RoboCup 2005.



A game scene during RoboCup 2005 in Osaka, Japan. The FU-Fighters robots have the blue markers.

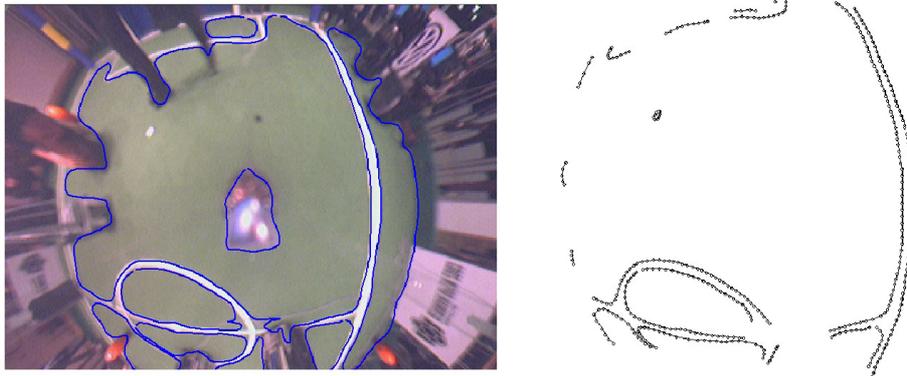
Around 12 students worked under my supervision each year on this project, which was funded by the Deutsche Forschungsgemeinschaft (DFG) from 2002 to 2008. Dozens of FU students took part in the eight world championships in which the FU Fighters participated.

R2: Adaptive Vision for Mobile Robots

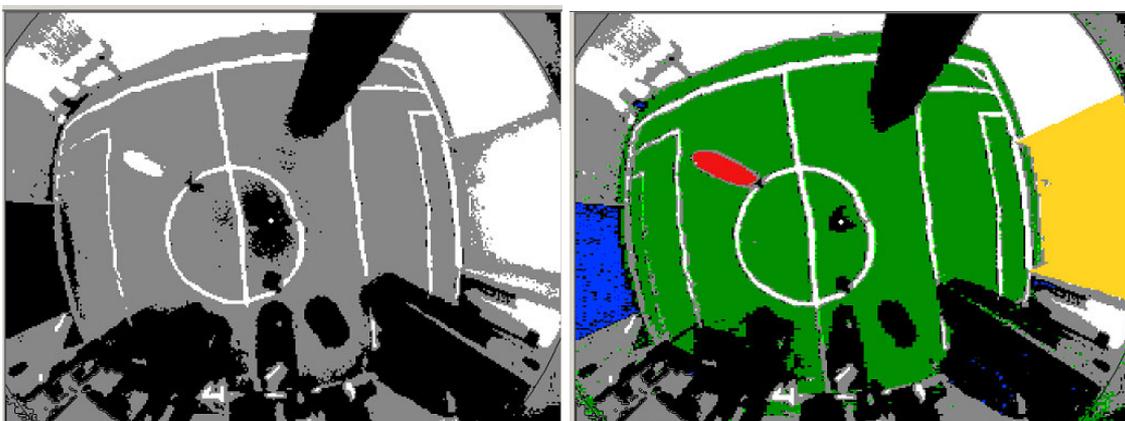
The FU-Fighters medium-size robots are autonomous and fast. The environment for these robots is an 8 meters by 10 meters large flat green field with white lines. The robots find their positions on the field by processing 30 frames per second of a video image. The laptops in the robots segment the image and find green regions on the field. The white lines are recognized and provide information about the position of the robots. Robust and fast localization techniques for mobile robots have been developed as part of this project. All of them rely exclusively on computer vision. No laser scanners or other types of sensors are used. The idea is to limit the robots to the same kind of sensors used by humans.

In our project we have developed new algorithms for a) fast and efficient segmentation of color images, b) real-time recognition of features on the field, and c) real time localization of the robots. Localization of the robot is done using Expectation Maximization, Kalman filtering, and Markov localization methods.

The colors recognized by the robots are not adjusted by human operators. They are learned by the robots. The robots initially self-localize on the field by using only black and white information (the white lines are still recognizable). From an initial localization, the robot formulates hypotheses about the colors of the regions enclosed by white lines and then computes its position again, possibly by displacing or rotating. Then new hypotheses are formulated, and so on. The whole color table can be filled in one or two seconds using this approach. The table can also be updated while playing.



Real-time segmentation of video images and contour extraction

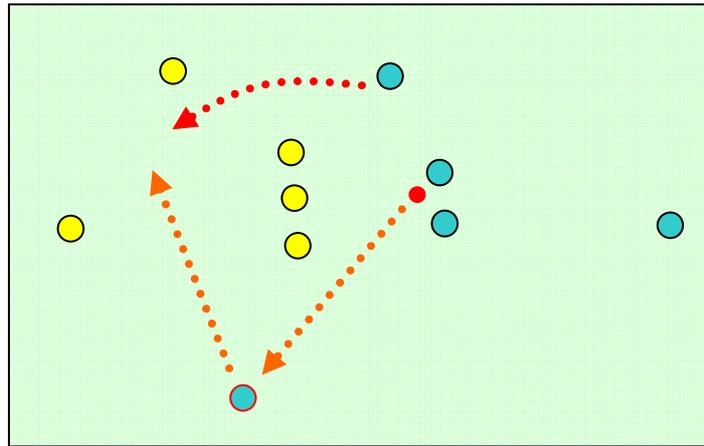


Initial bw segmentation (left) automatically learned color calibration (right)

Dr. Felix von Hundelshausen (now at the University of the Bundeswehr) directed this project from 2001 to 2005. Ketill Gunnarson (now at Google) took over from 2005 to 2007.

R3: Neural Coach for Autonomous Robots

We have been programming autonomous mobile robots for robotic competitions. Up until now, the robots' behavior was programmed directly in great detail. In this project we are investigating ways of controlling the robots with programs that can learn from experience. We want to teach the robots to make moves as if they were human players being instructed by a coach. We want to give them static examples of a play situation and we expect them to generalize to new, unseen, situations. The figure shows an example of the way a human coach would explain to a team of human players (playing five-a-side soccer) how to play to get around a barrier of opponents.



A coach would explain a move to human players using this kind of diagram

As can be seen from the diagram, the human coach uses an abstract form of communication, a sketch of the game situation, to illustrate the strategy needed to move two players forward. The human coach does not need to draw all possible variations of a barrier of opponents for the team members to understand the tactic: pass to the left, and then try to double-pass to the other player, who should have come forward. Using this approach, we do not have to program all game alternatives by hand. We can just compile a set of static examples of “good” and “bad” moves. We can then use machine learning to adapt the behavior software to the learned scenarios.

The static field situation depicted in the diagram has to be encoded as a set of features that can be processed by one or several neural networks. The network learns, for example, what a “good” passing situation is, and can also learn to show the robots the field positions with a high probability of reward. Using the output of such networks, the robots can anticipate game positions. They stop being purely reactive and become more forward-looking. It has been said that most of our own conscious processing goes into *predicting the future*, that is, formulating expectations. When we walk, and when we see, we predict our future position and even the position of our eyes. We know that we are moving in the world, the world is not moving around us. Predicting the best possible movements for the robots amounts to making a strategic evaluation of the alternatives the opponent has. This is exactly what we want to achieve.

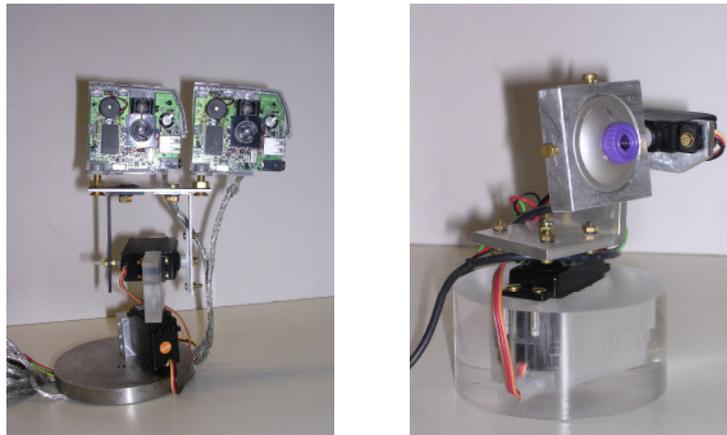
The first version of the neural coach was used during RoboCup 2005. The FU-Fighters won first place at that competition.

This project was pursued by Bastian Hecht (now a PhD student) under my supervision.

R4: Stereoscopic Head for A Biped Robot

We would like mobile robots to be able to navigate in human industrial or office environments using only visual information. In this project, we developed a robotic head able to find and track objects. The robotic head consists of two cameras providing two images of the environment. A laptop processes both images and computes the distance from the head to segmented parts of the image. The robotic head can recognize features found in human faces and can thus fixate a person's face. If the person moves, the cameras track the person's face. The cameras have three coupled 3D degrees of liberty. Control is done by PICs, which deliver signals to servomotors. The challenge in this project is to develop real-time stereoscopic processing algorithms for recognition and tracking of objects.

The robotic head was enhanced with microphones for speech processing and has been in several robots. A new version is now being used in our humanoid robots (FUmoids) and will be integrated into a human size humanoid in 2010.



Two cameras with servo motors capable of rotating around two axes.

This project was pursued by Erik Cuevas and Daniel Zaldivar (postdocs, now professors at the University of Guadalajara) under my supervision and with funding from Deutscher Akademischer Austauschdienst (DAAD).

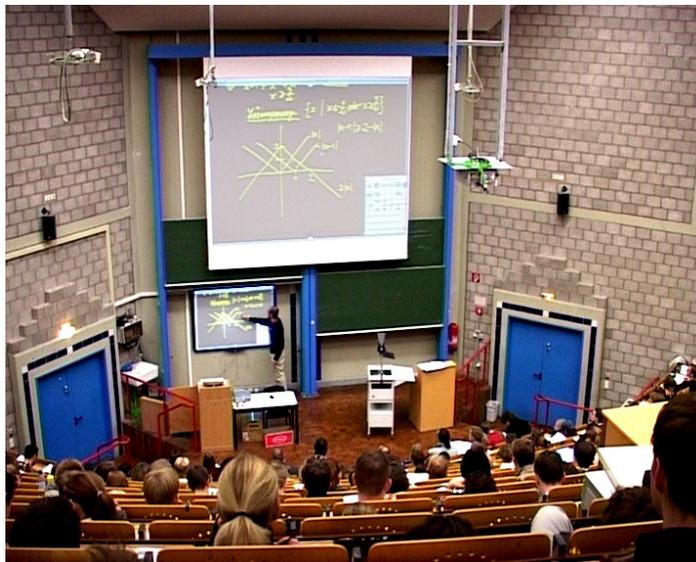
Human-Computer Interaction and E-Learning

E1: E-Chalk – An Intelligent Chalkboard

www.echalk.de

E-Chalk is software we developed in order to enhance the classroom experience. Our guiding philosophy was to provide the teacher with a tool to make lectures more dynamic and interesting. At the same time, the lectures can be stored and transmitted live over the Internet. E-Chalk transforms a large touch-sensitive screen into a smart teaching tool. The instructor writes on the screen using a special stylus: the color and width of the line strokes can be changed so the screen emulates a classical chalkboard. However, the lecturer can do much more with E-Chalk: images from bookmarks can be pasted to the screen, a query to a remote Web service can be sent, an algebraic server can be asked to plot a function, Java applets with animations can be pasted to the board. A copy of the lecture (audio, board, and video) is stored on a server and the lecture can be also transmitted live over the Internet. A remote viewer needs only a conventional Java-capable browser in order to view or replay a lecture. A PDF transcript of the lecture is available from an automatic repository.

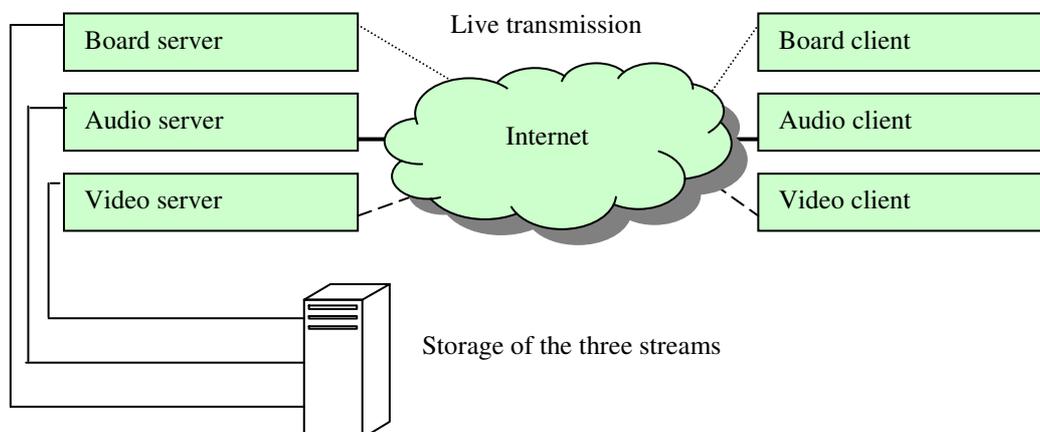
The E-Chalk architecture is based on the metaphor of the classical chalkboard, enhanced by intelligent assistants running in the background, and providing various services. One assistant is responsible for interpreting the handwritten input of the user. Another is a mathematical formula recognizer that processes handwritten queries for the algebraic server. A circuit simulator recognizes sketches of digital circuits and runs a simulation. An algorithm simulator accepts sketches of graphs as input data and runs graph algorithms, animating them on the screen. Additional intelligent assistants can be interfaced by E-Chalk savvy users by just interfacing their applications to E-Chalk using the API we provide.



A professor teaching with E-Chalk at the Technical University Berlin. The lecturer writes directly on the contact-sensitive whiteboard, a large copy of the image is projected above him. More than 1500 students were taught mathematics for engineers with E-Chalk in 2003.

The E-Chalk project was funded by the Deutsche Forschungsgemeinschaft (DFG) and the Klaus-Tschira-Foundation. Several students work on the project under my supervision. Dr. Lars Knipping (now TU Berlin) and Dr. Gerald Friedland (now in Berkeley) obtained their PhD degrees with this project.

E2: Software Architecture for a Smart Teaching Board



In this project the board components of E-Chalk and an interface for third-party applications are being developed. Intelligent software agents look over the shoulder of the lecturer and handle various types of input. In 2007, we ported the system to MPEG4 so that E-Chalk lectures can now be played on computers, cellular telephones, and iPods.

The E-Chalk server software currently consists of three main and several minor components. The three main components are the audio, the video, and the board server. The minor components are the database manager and the PDF converter. The E-Chalk client system consists of a set of independent receiver applets that synchronize themselves by communicating through our Media Applet Synchronization Interface (MASI).

E-Chalk has basically two ways of operating: online for live transmissions and offline for archived transmissions. During live transmissions the audio stream, video stream, and board events are recorded from their devices, compressed, and sent. They are then received, decompressed and replayed in real time. While this is being done, everything that is sent out is also stored in files. As a result, the E-Chalk client system has two modes. In live mode, each client connects to its corresponding server through a socket connection. In on-demand mode, clients use an http connection to receive the files and no E-Chalk server is needed.

We are working on interfacing more third-party programs. Mathematica, Maple, and other programs can already process input coming from the board. Algorithmic animations of handwritten input are now possible. We also have a circuit simulator that can evaluate logical circuits drawn with AND, OR, and NOT gates. All these applications rely on the E-Chalk API we have developed. The API allows a program to handle handwritten input from E-Chalk and to request strokes to be drawn on the board. We call this program “intelligent assistants”.

This project was pursued by Lars Knipping (postdoc) and Philip Holzschneider under my supervision, with funding from the BMBF (Minister for Education and Research).

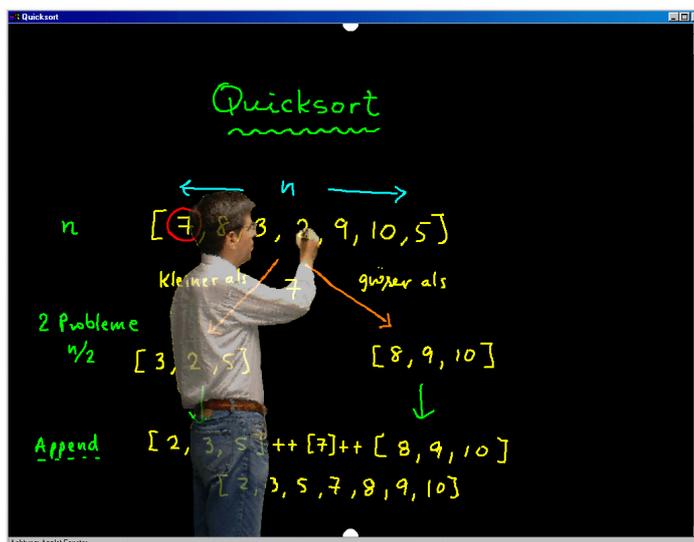
E3: Intelligent Audio and Video Processing for E-Chalk

In this project we developed software to increase the quality of the audio channel and to merge the board and video channels in the future.

The quality of the audio channel suffers during lecture hall recordings for several reasons (noise, echo, reverberance, etc.). It can be improved by using special hardware and a mixer. But in our case, we want to provide the user a way of producing high-quality recordings purely by adapting the software. Our software is an expert system for audio recordings that extracts a fingerprint of the noise from the hardware, as well as a fingerprint of the recording situation. The fingerprint is used to eliminate artifacts and to modulate the recording level. The results can compete with expensive hardware.

In the second part of the project we want to track and segment the lecturer from the video channel. The segmented lecturer can then be transmitted together with the board data and can be superimposed on the board. The effect is that the board image is reproduced in high quality (as board strokes) and the lecturer's facial expression is still visible. The viewer can even dim the lecturer in order to see through him.

In this project we are extending state of the art algorithms for lecturer recognition and tracking. We need to be able to cut the lecturer image in real-time. Reconstruction at the receiving end must also be done in real time. A first version of the lecturer segmenter has been written: it can segment the lecturer at 10 frames per second. A version for tracking objects in screenplays is being written by an open source group affiliated to GIMP.



An image of future results: the lecturer video is pasted onto the whiteboard and can be dimmed

This project was pursued by Gerald Friedland under my supervision. We collaborate with the GIMP open source project.

E4: Recognition of Handwritten Mathematical Formulas

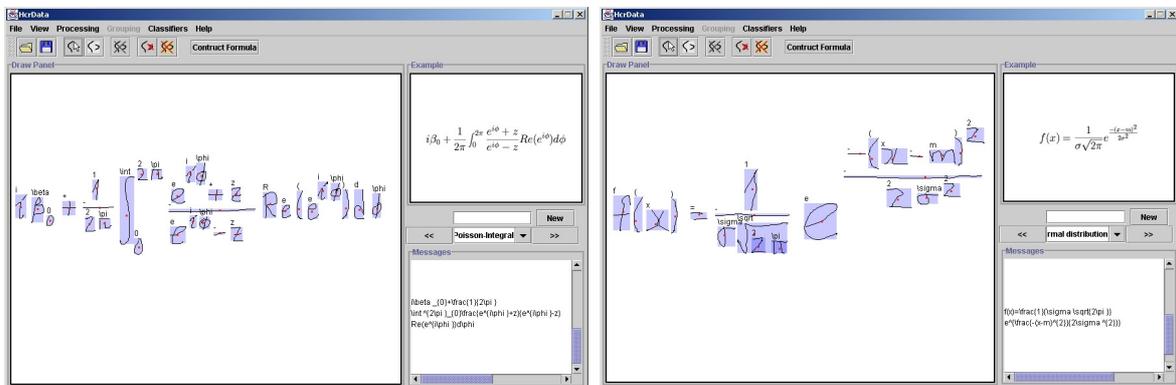
<http://mathfor.mi.fu-berlin.de/>

In this project we are developing a mathematical formula recognizer for E-Chalk. The idea is to give the lecturer a tool for evaluating formulas on-the-fly during a class. The formula recognizer can also be used to transform handwritten formulas into LaTeX code.

When a formula is drawn free-hand on the board, the recognizer first identifies and groups individual strokes. Once the strokes belonging to a character have been selected, the next step is to simplify their shape, which is done by eliminating intermediate points in sequences along lines. Only some points are kept in curves, i.e. those which can best approximate the curve piecewise within a certain tolerance factor.

When a character has been reduced to simplified strokes, represented by just a few points each, we encode each stroke as a vector of features. Once each stroke has been encoded as a vector, we train classifiers for each character (taken from a training set produced by several people). The classifiers are neural networks or support vector machines. There is a classifier for each different number of strokes (since the number of encoding vectors is different). The classifiers recognize each letter and provide their ASCII code.

Finally the two-dimensional layout of the formula must be recognized. We do this using the minimum spanning tree of the symbols, modified by some appropriate heuristics. The formula recognizer is currently in use in E-Chalk and JMath, the formula recognition shell, can be used to produce LaTeX code for handwritten formulas.



Examples of mathematical formula recognition. The LaTeX code is shown on the respective lower right subframes, the typeset formula on the respective upper right subframes.

This project was pursued by Ernesto Tapia (postdoc) under my supervision and with funding from DFG (Deutsche Forschungsgemeinschaft)

E5: Algorithmic Animation for E-Chalk

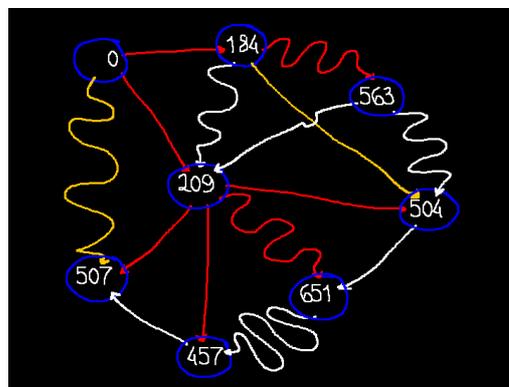
Algorithmic animation is the art of producing a visual image of the operation of a certain algorithm. Several algorithmic animation systems have been written, but none of them has ever handled handwritten input. In this project we will write the pseudocode for an algorithm on the electronic chalkboard. The computer will then recognize the handwriting of the user, will compile the program and semi-automatically produce the required animation. In the figure below, for example, the user has written a program. Encircling the variables A and i tells the system that we want to run the program seeing a visual representation of A and i . The array A is shown as a collection of bars. The variable i is shown as an arrow, because it is an index to the array A .

```
A = [10,9,8,7,6,5,4,3,2,1]

for j in range(1,10):
    key = A[j]
    i = j - 1
    while( i >= 0 ) and (A[i] > key):
        A[i+1] = A[i]
        i = i - 1
    A[i+1] = key
```

The main idea: handwritten programs are executed and selected variables are visualized

It is already possible to run algorithms in E-Chalk, for which the input is entered as a sketch. Dijkstra's shortest path algorithm can be programmed and made to run on a graph entered by the user by hand. The figure below shows an example.



An animation of Dijkstra's shortest path algorithm running with hand-drawn input.

This project was already the subject of a PhD thesis and will be continued using Python as the algorithmic language.

E6: Hardware Design for a Multimedia Classroom

At the Free University Berlin we designed and built a novel multimedia classroom. The picture below shows the electronic blackboard used for courses (6 meters long). The screen is illuminated by four LCD rear projectors, emulating a chalkboard. The lecturer can write directly on the board using a light pen. The light of the pen is recognized by four video cameras located behind the screen. The video cameras find the red light blob in the image and provide its coordinates, through a driver, to any program requiring mouse input. Our E-Chalk software can be operated in this manner. The light pen is a cheaper alternative to laser scanners, magnetic grids, and other existing methods of digitizing the position of a stylus.

We also built a special stylus connected by Bluetooth to the main computer. The stylus contains 20 buttons, which can be pressed by the user. They allow the lecturer to change color, line width, or scroll the chalkboard without having to access the menu on the screen. This is an efficient alternative for large computer screens, such as the one shown in the picture.

In this project, we are investigating different methods of digitizing the stylus position, using ultrasound or light. We have developed automatic calibration methods for such systems. We are also investigating the design of human computer interfaces for large screens.



The new electronic classroom at the FU Berlin

In 2008 we incorporated a new electronic voting system that we developed. It is cheap and can manage up to 200 students.

The project was funded by the FU Berlin as one of its E-Learning initiatives. Christian Zick (technical staff) lead the project under my supervision.

History of Computing

H1: Konrad Zuse Internet Archive www.zib.de/zuse

Konrad Zuse is considered the father of the computer in Germany. We have been collecting Zuse's papers and diagrams in a systematic way since 1998. Some of the documents were typewritten and others were scanned in PDF format. The result is the Konrad Zuse Internet Archive, present on the Internet since 1999. The archive contains most of the important documents and circuit diagrams produced by Zuse and his company between 1936 and 1995. The archive also contains simulations of some of the machines, especially the Z3, which was the first programmable computer in the world when it was unveiled in 1941. The simulations were written in Java.

The archive also contains two important pieces of software. One of them is an editor and interpreter for Plankalkül, the high-level programming language designed by Zuse in 1945. No compiler or interpreter was available until 2000 when our system was put on the network. The other is our implementation of Konrad Zuse's 1945 chess program in Java. It was the first real chess program and the interested user can experiment with it.

The archive contains pictures of our hardware reconstructions, most notably the Z3 and its addition unit. The number of personal visits to the archive has reached around 100,000 per year. The Internet visitors are from Germany and several other countries.



The opening page of the Konrad Zuse Internet Archive

The project was funded by the DFG and Konrad Zuse Zentrum Berlin.

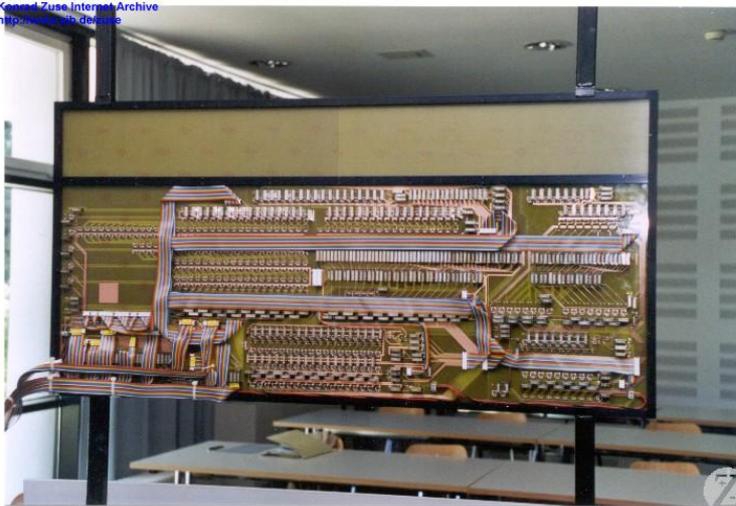
H2: Reconstruction of Konrad Zuse's Z3 Computer

We built a modern reconstruction of Konrad Zuse's Z3 computer, using electromechanical relays for all the logical components. In 2008 we will be working on a new version of the addition unit printed with Plastronic components.

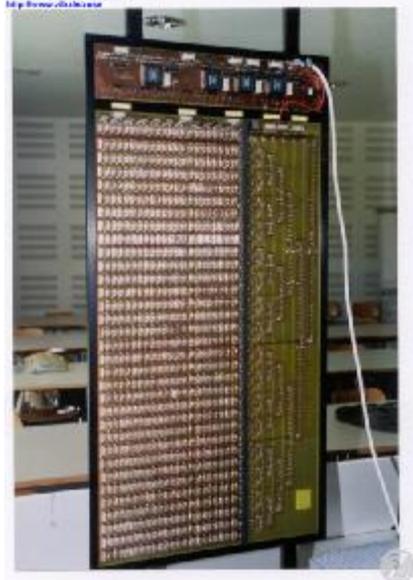
The reconstruction started in 1997, when I first deciphered the construction plans of the Z3. I published the result of this investigation in 1998 (*Die Rechenmaschinen von Konrad Zuse*, Springer-Verlag). Shortly afterwards, we started working on the reconstruction of the Z3's addition unit. Together with Dr. Frank Darius (TU Berlin) and Georg Heyne (Max-Planck-Gesellschaft) we produced ten addition units, which were distributed to several German and American universities. The next step was to produce a working copy of the whole machine. The same team worked from 2000 to 2003 until the machine was finished. It is now part of the collection in the Konrad Zuse Museum in Hünfeld, Germany.

The main challenge during reconstruction was to use modern components while retaining the original circuits of the machine. We opted for simulating the console of the machine with a computer, in order to eliminate purely mechanical components and extend the working life of the machine. The pictures below show some views of the Z3 reconstruction, a project which involved ten people from different institutions over a period of three years.

Konrad-Zuse-Internet-Archive
<http://www.zuse.de/usa>



Konrad-Zuse-Internet-Archive
<http://www.zuse.de/usa>



View of the back of the processor (left) and memory (right) of the Z3

The project was funded by the DFG, Gesellschaft für Informatik, Konrad Zuse Zentrum Berlin, and private donations from Horst Zuse and his wife.

H3: Simulation and Reconstruction of the ENIAC

The Electrical Numerical Integrator and Computer (ENIAC) was built at the Moore School of Engineering, University of Pennsylvania, and has been hailed as the first large-scale electronic computer in the world. In this project we wrote the first software simulation of the ENIAC. The simulation is written in Java and runs as a Java applet on any computer connected to the Internet.

The main difficulty encountered in simulating the ENIAC is that it did not use software. The program was hardwired by special programming assistants. Figure 1 below shows the cycling unit of the ENIAC, the initiating unit, and two accumulators. The user can activate the switches and connect accumulators with transmission lines for digits or pulses. Once the program has been “cabled”, the machine can run at the velocity desired by the user.

This simulation provides computer science students with a better understanding of the architecture and programming of the ENIAC. The simulation was shown in 2004 in Germany and the US, and is currently online at www.zib.de/zuse.

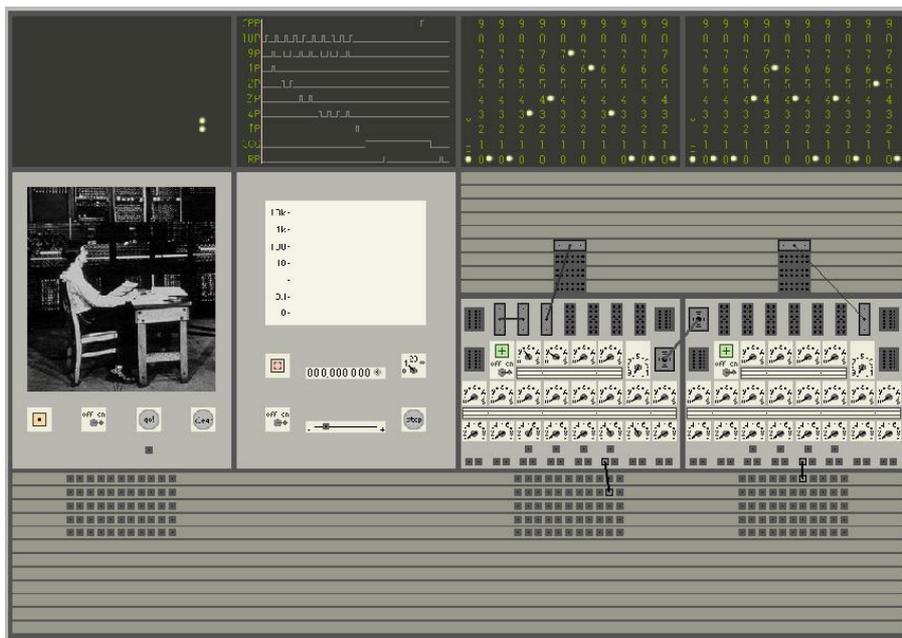


Figure 1: Screenshot of a portion of the ENIAC simulation showing two accumulators.

In the simulation of the ENIAC, up to 20 accumulators, two constants units, and other units can be assembled. Appropriate signal lines allow the user to perform tabular computations, iterations, and even to insert “if-then” commands. Since the machine is so large, the simulation allows the user to zoom in or out of the machine to get a better impression of the overall cabling. The figures below show some screenshots of the simulator.

This project was pursued by Till Zoppke under my supervision.