AI & Robotics @ Edinburgh six decades … and beyond

Austin Tate

Emeritus Professor of Knowledge-Based Systems

AIAI, University of Edinburgh

The New Club, Edinburgh, 6th January 2022

Artificial Intelligence @ Edinburgh over the Years

- Experimental Programming Unit (EPU), 1963-1966
- Department of Machine Intelligence and Perception (DMIP), 1966-1970
- Department of Machine Intelligence, 1970-1973
- School of Artificial Intelligence, 1973-1974
  + Machine Intelligence Research Unit (MIRU), 1973-1977
- Department of Artificial Intelligence (DAI), 1974-1998
  + Artificial Intelligence Applications Institute (AIAI), 1983-2019
- In 1998, the University joined together three departments: Artificial Intelligence, Cognitive Science and Computer Science, as well as a number of research institutes including AIAI and the Human Communication Research Centre, to form the School of Informatics.
- Institute for Representation and Reasoning (IRR), School of Informatics, 1998-2001
- Centre for Intelligent Systems and their Applications (CISA), School of Informatics, 2001-2019
- Artificial Intelligence and its Applications Institute (AIAI), School of Informatics, 2019-

A number of other departments and schools at the University of Edinburgh as well as other research institutes in the School of Informatics work on a range of topics within the field of Artificial Intelligence.

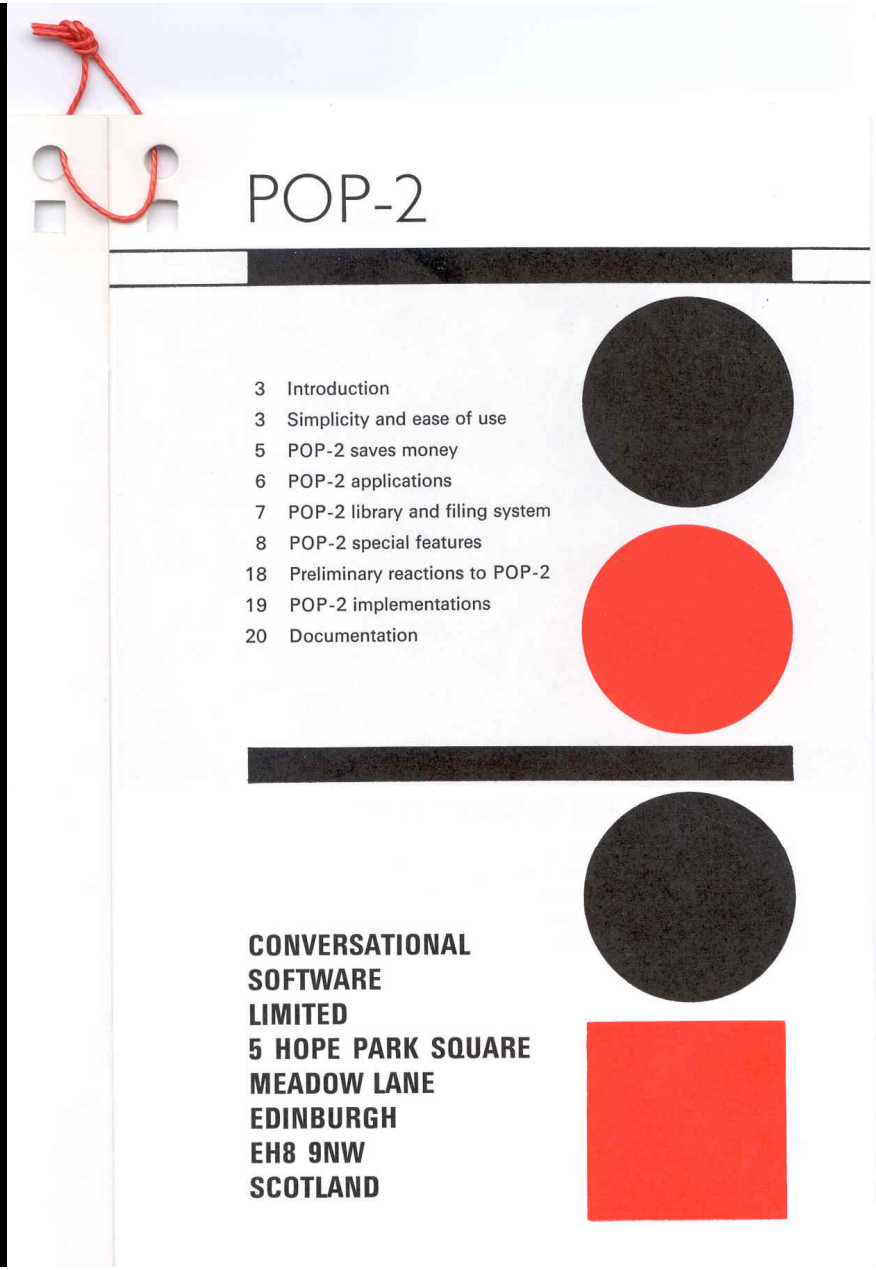

First AI Company, 1969

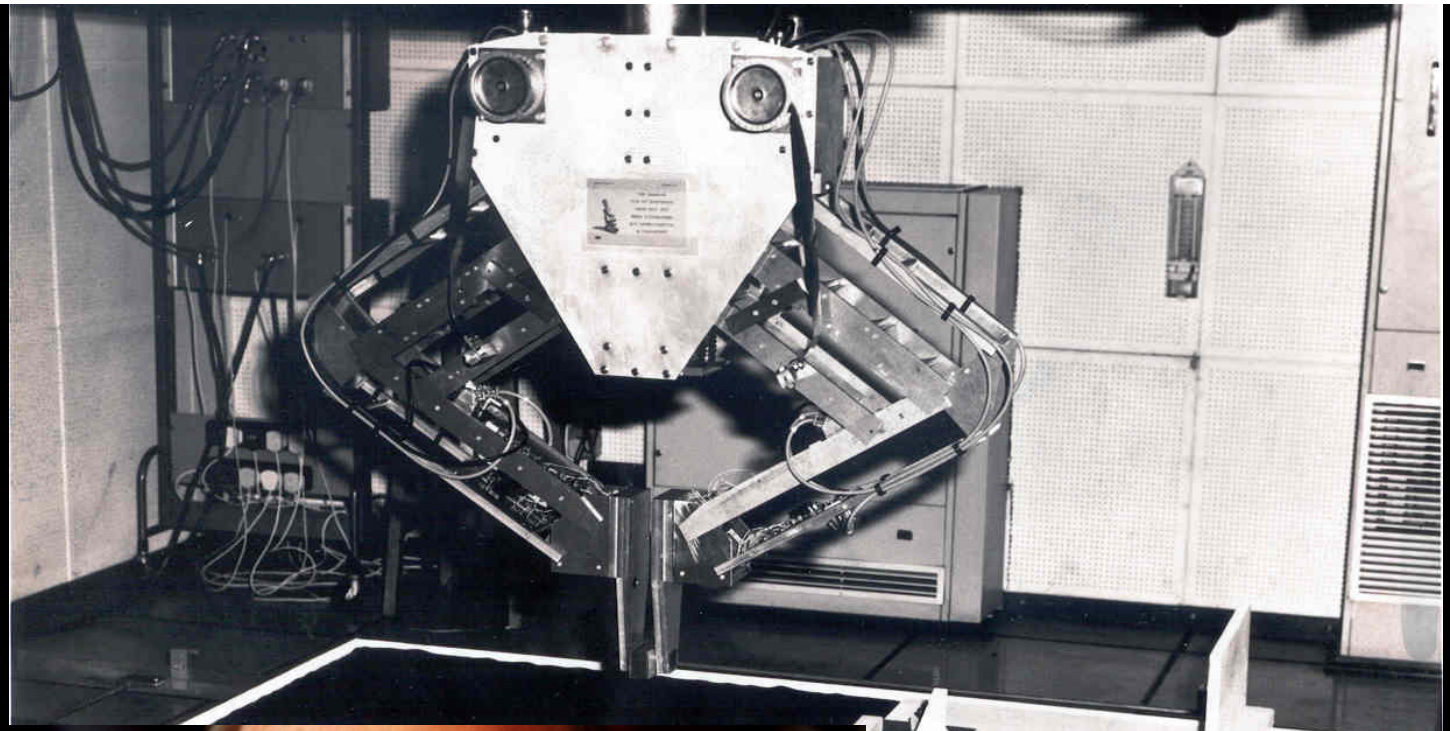

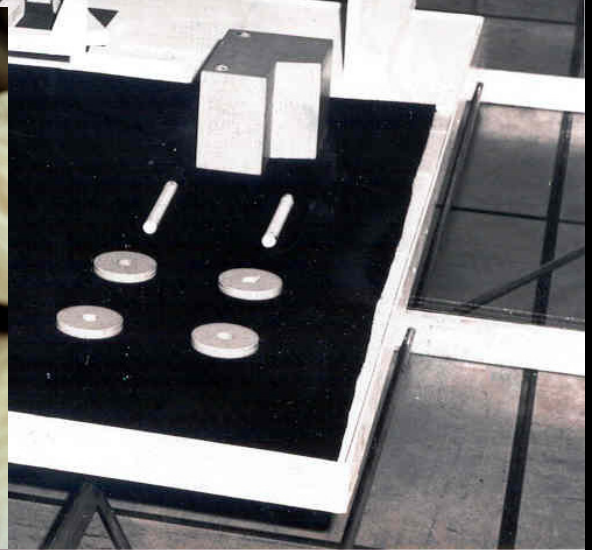

Freddy II, 1972


ALVEY PROGRAMME
ANNUAL REPORT
1984
ALVEY
ALVEY PROGRAMME
ANNUAL REPORT
1985
ALVEY
ALVEY PROGRAMME
ANNUAL REPORT
1985
POSTER SUPPLEMENT
ALVEY
CONFERE
1986
ALVEY
1987
INDUSTRY
MOD
DTI
SERC
ALVEY
PROGRAMME
1987
ALVEY PROGRAMME
ANNUAL REPORT
1988
ALVEY

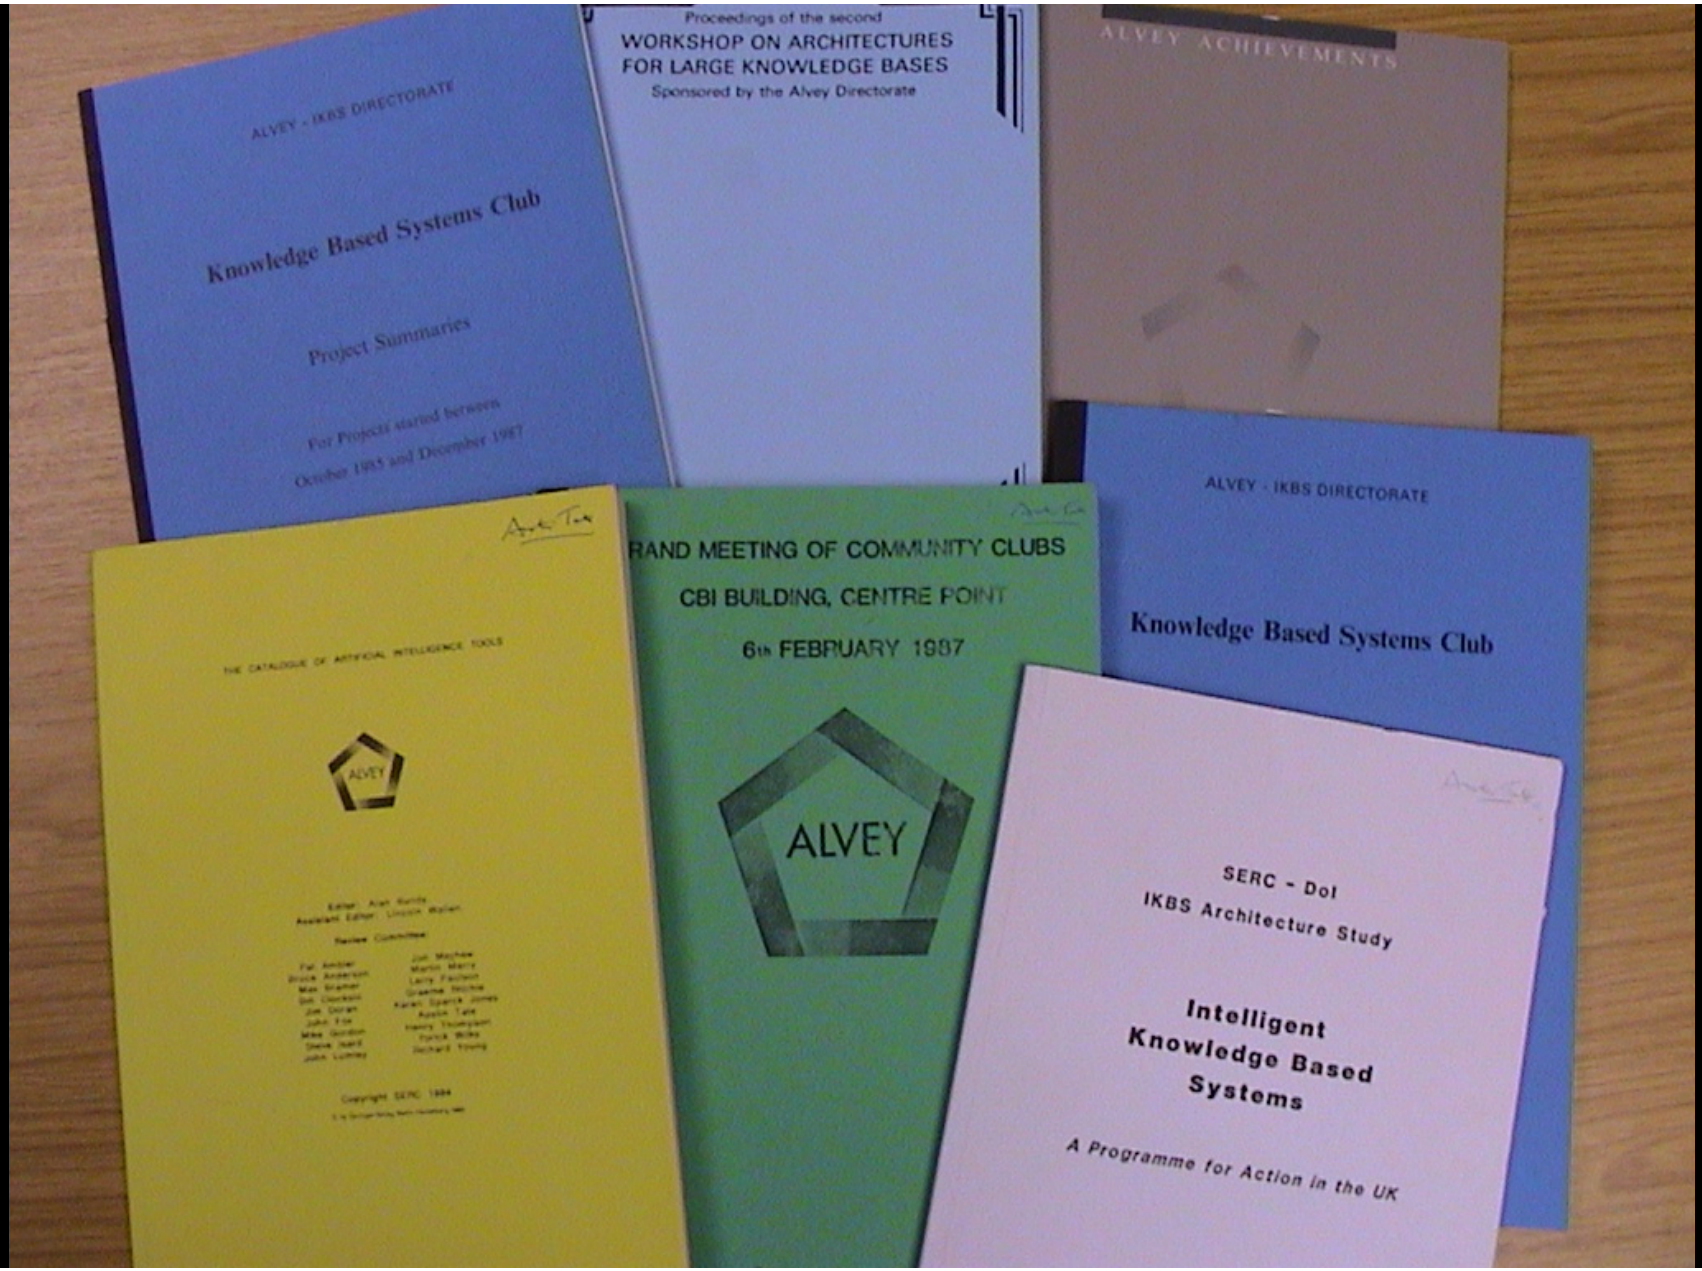
ALVEY - IKBS DIRECTORATE
Knowledge Based Systems Club
Project Summaries
For Projects started between
October 1985 and December 1987
Proceedings of the second
WORKSHOP ON ARCHITECTURES
FOR LARGE KNOWLEDGE BASES
Sponsored by the Alvey Directorate
ALVEY ACHIEVEMENTS
ALVEY - IKBS DIRECTORATE
Knowledge Based Systems Club
THE CATALOGUE OF ARTIFICIAL INTELLIGENCE TOOLS
ALVEY
RAND MEETING OF COMMUNITY CLUBS
CBI BUILDING, CENTRE POINT
6th FEBRUARY 1987
ALVEY
SERC - DoI
IKBS Architecture Study
Intelligent
Knowledge Based
Systems
A Programme for Action in the UK

The Establishment of an

ARTIFICIAL INTELLIGENCE APPLICATIONS INSTITUTE

at the
University of Edinburgh

University of Edinburgh

A · I · A · I

VERSITY OF EDINBURGH

inburgh University started work in Artificia
as the only Department of Artificial Intellige
esult, it now hosts the largest and most p
with over 200 people researching in all as
of Artificial Intelligence, Computer Sci
Edinburgh, form one of the strongest
centres in Britain.

ARTIFICIAL INTELLIGENC

The techniques and tools develo
reached the point where they c
commercial and industrial relevan
response to this industrial intere

AIAI is a technology transfer org
interpret and present resea
knowledge based systems for
academia.

AIAI is located in central
Department of Artific
Technology Research.

TECHNICAL G

AIAI is organised
specialist. The gr

- Knowled
- AI Pro
- AI Pa

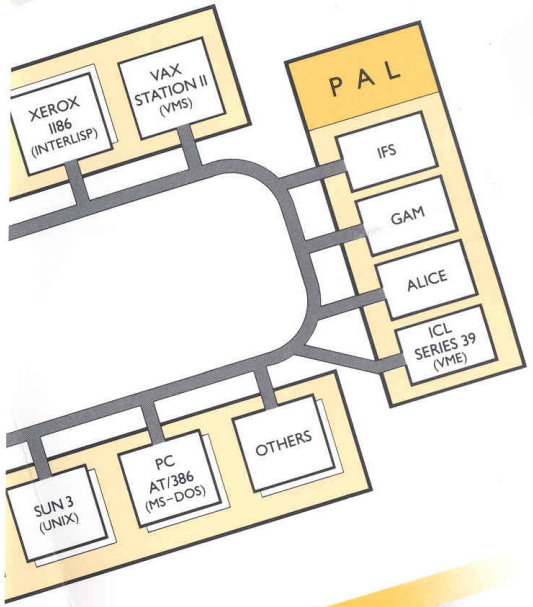

PAL

AIAI hosts a national Parallel Architectures Laboratory (PAL) which aims to encourage cooperative work on novel architectures and their applications within AI. A variety of these are currently being investigated, including software systems such as Parallel Prologs and hardware systems such as the ALICE graph reduction engine and the FACT/GAM (Generic Associative Memory) and IFS (Intelligent File Store), developed under the Alvey programme.

Further architectures are accessible through the University's Parallel Computin Consortium of which AIAI is a member. PAL is jointly supported by DTI, SERC a Alvey.

established
ardware and
er substantial
and suitability
oolkits available

)PLOG
)PSS

ble to assist visiting
ne Alvey Programme,

AIAI, 1984

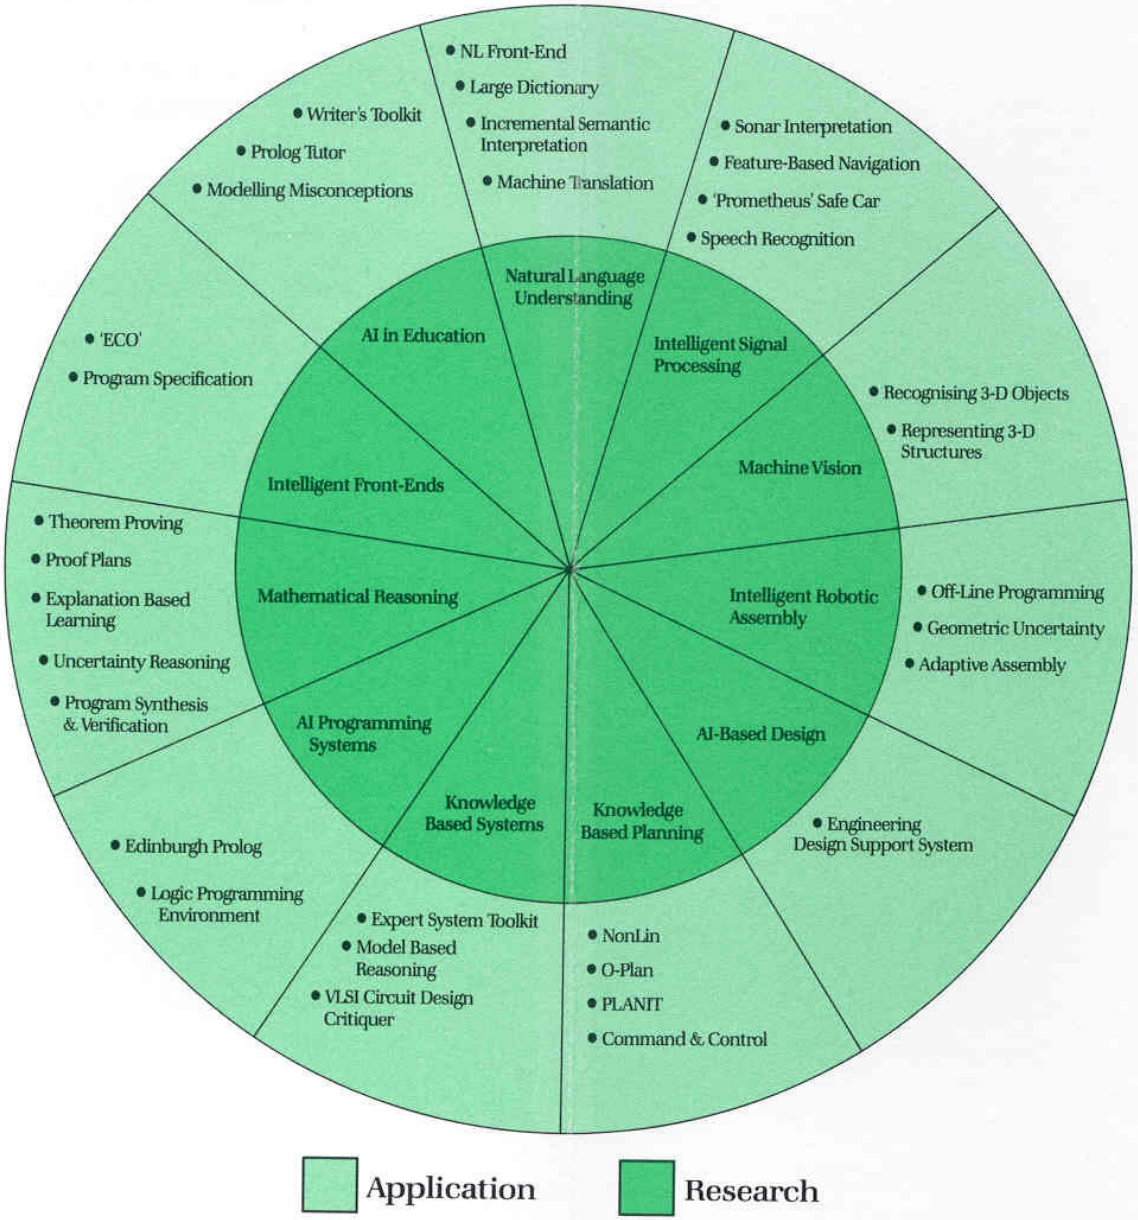

Research Topics, 1980s

Call for participation

The scientific scope of the conference includes, but is not limited to the following topics:

List of area keywords

abduction
adaptive systems
AI architectures
AI and data integration
AI and creativity
AI and security
AI and the Internet
applications
art and music
artificial life
automated modeling
automated reasoning
autonomous agents
belief revision and update
business intelligence
case-based reasoning
causality
cognitive modeling
cognitive robotics
common-sense reasoning
computational complexity
computer-aided education
conceptual graphs
configuration
constraint programming
constraint satisfaction
corporate knowledge
customer relationship management
data mining
decision theory
decision trees
description logics
design and configuration
dialog processing
diagnosis
discourse modeling
distributed AI
enabling technologies
expert systems
game playing
genetic algorithms
geometric reasoning
heuristics
human computer interaction
inductive logic programming
information agents
information extraction
information retrieval
intelligent databases
intelligent e-mail processing
intelligent query processing
knowledge acquisition
knowledge discovery
knowledge engineering
knowledge representation
lifelike characters
logic programming
machine learning
machine translation
market-oriented programming
mathematical foundations
model-based reasoning
multiagent systems
multimedia
natural language
negotiation
neural networks
non-classical computation models
nonmonotonic reasoning
ontologies
perception
personalization
planning
philosophical foundations
predictions
probabilistic reasoning
problem solving
qualitative reasoning
reactive control
real-time systems
reasoning about actions and change
reinforcement learning
resource-bounded reasoning
robotics
rule learning
satisfiability
scheduling
scientific discovery
search
semantic web
simulation
software agents
spatial reasoning
spatial and temporal databases
speech processing
temporal reasoning
theorem proving
uncertainty
user interfaces
user modeling
validation and verification
virtual reality
vision
web agents
web intelligence
web mining
web search

Typical AI Conference Topics

IJCAI 2003, Mexico


THE UNIVERSITY of EDINBURGH
INFORMATICS FORUM

Informatics Research Landscape

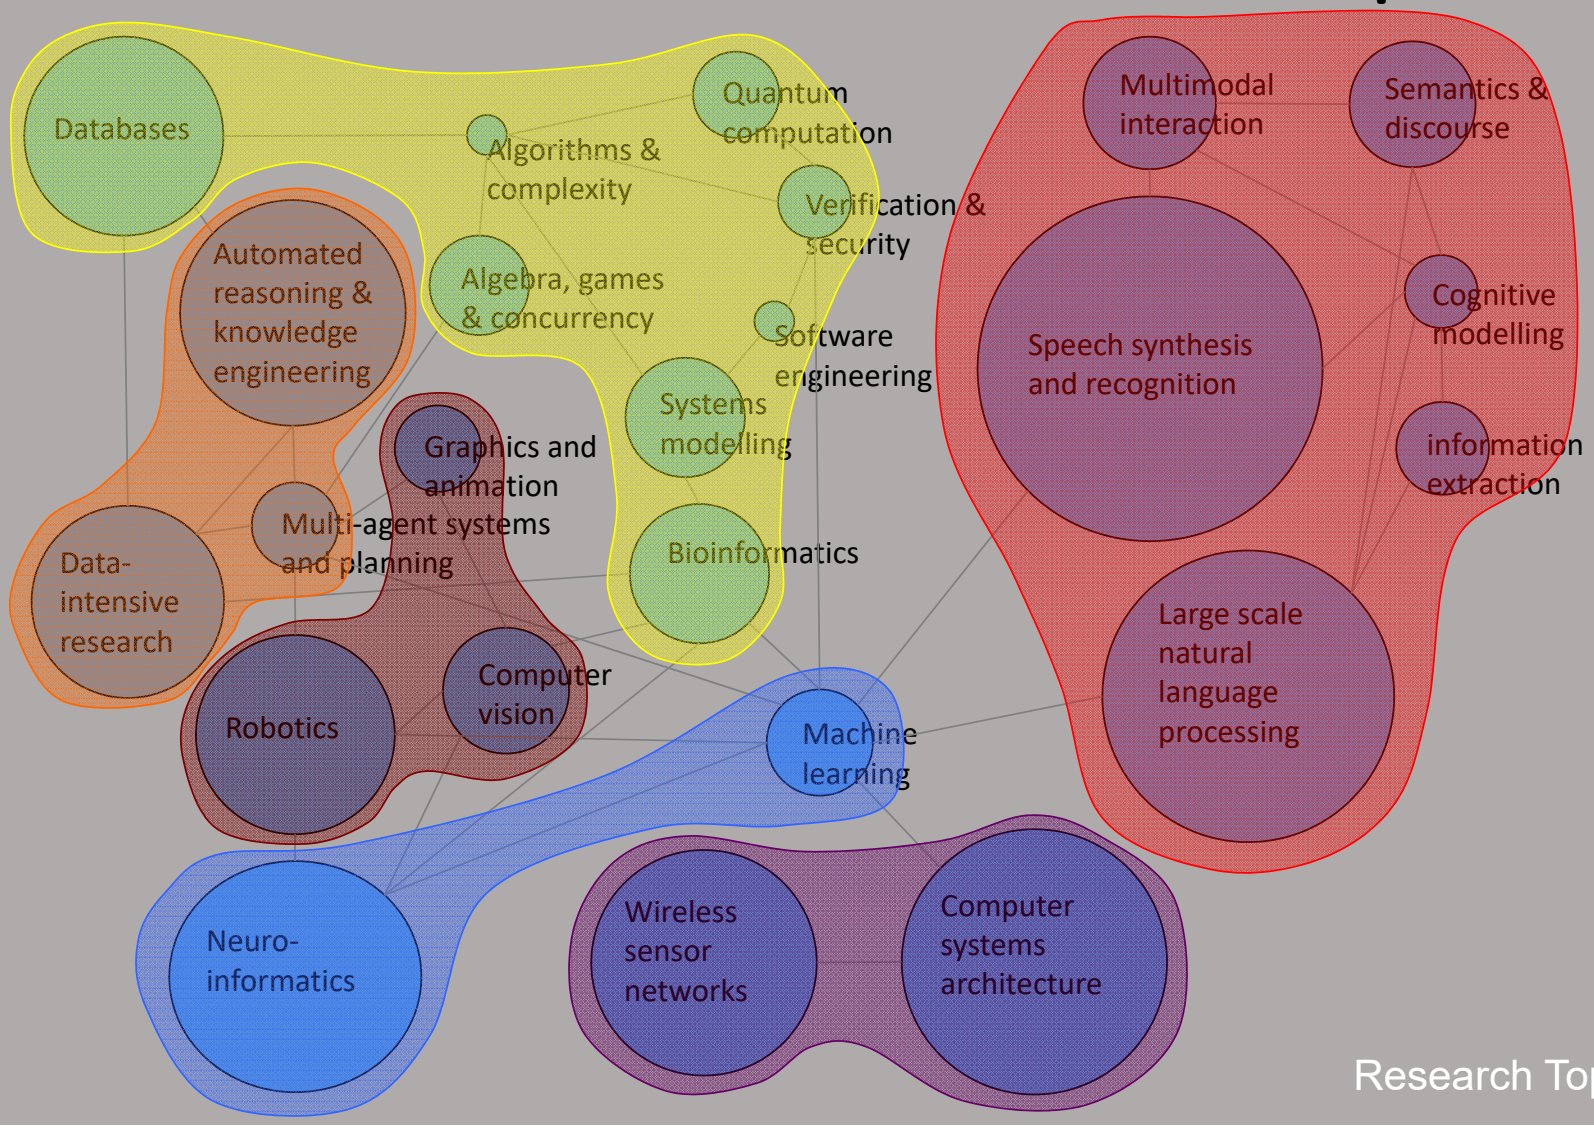

Research Topics, 1998

Informatics Institutes and Relationships

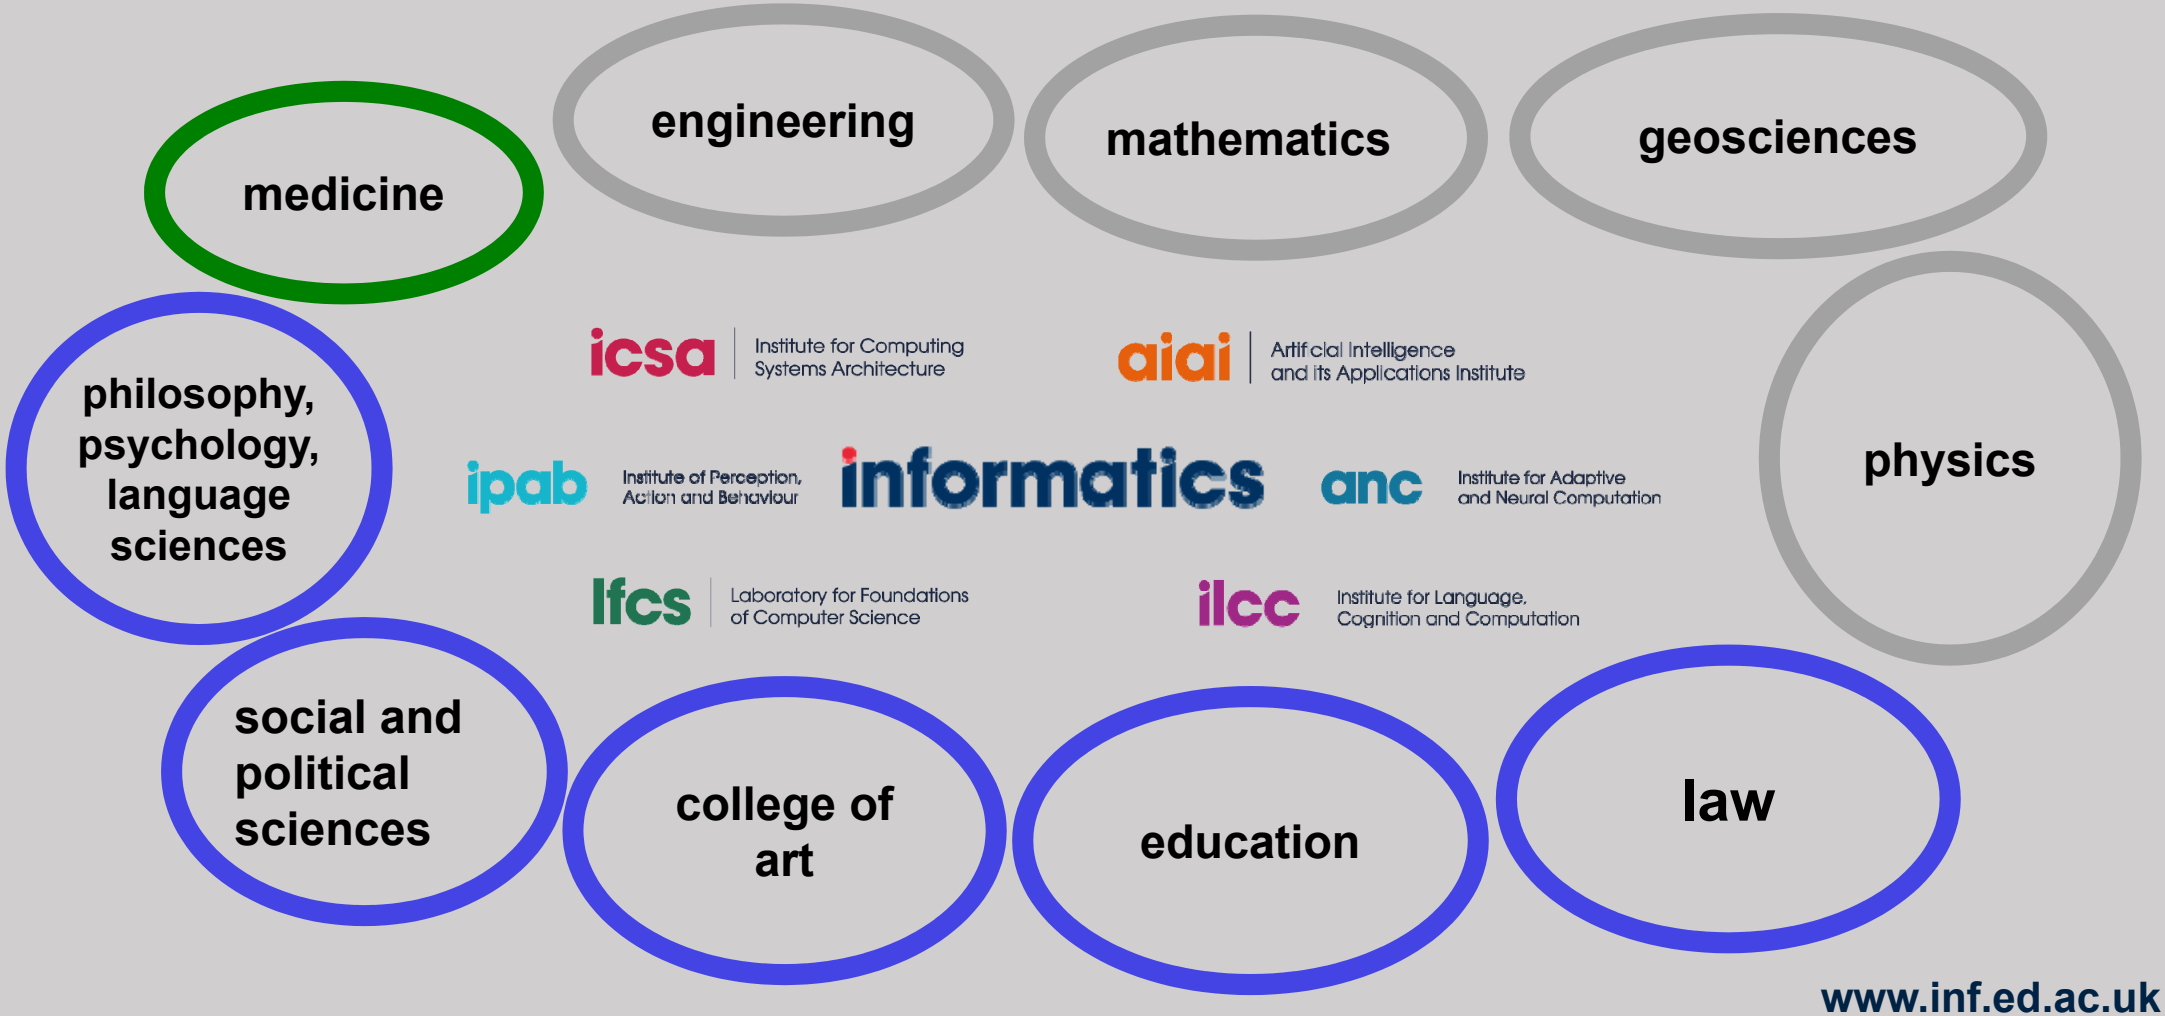

www.inf.ed.ac.uk

Edinburgh Startup Ecosystem

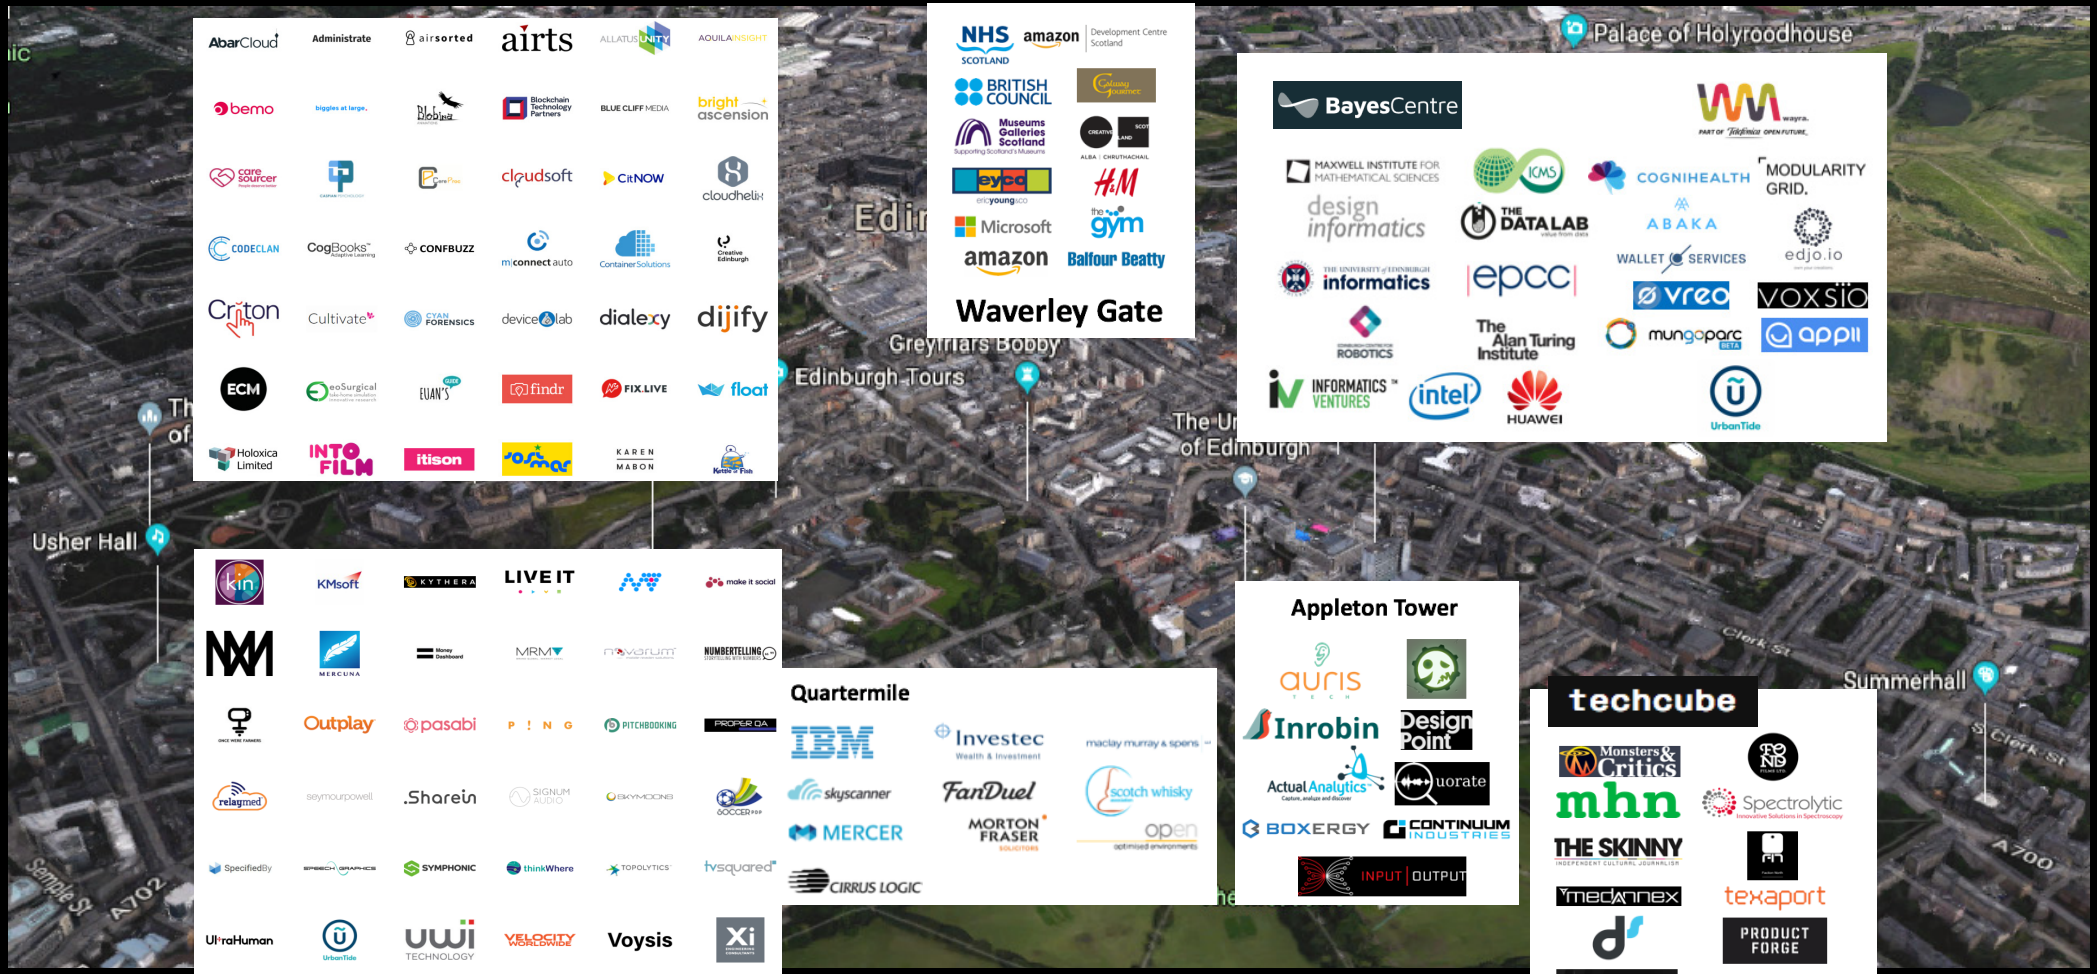

Bayes Centre
Making Data Matter

- Innovation Centre for Data Science and AI
- Putting Data Driven Technologies to Practical Use
- Research, Education, and Innovation

Bayes and the Bayes Theorem

Bayes Theorem - describes the probability of an event, based on prior knowledge of conditions that might be related to the event.

Underlies much of the “deep learning” data driven (“big data”) systems now being studied and applied in many practical situations.

Rev. Thomas Bayes, minister, philosopher & statistician 1701-1761, University of Edinburgh (logic & theology).


E-Commerce
Navigation
Robotics
Human Resources
Healthcare
Agriculture
Gaming
Vehicles
Social Media
Marketing

The Bayes Centre

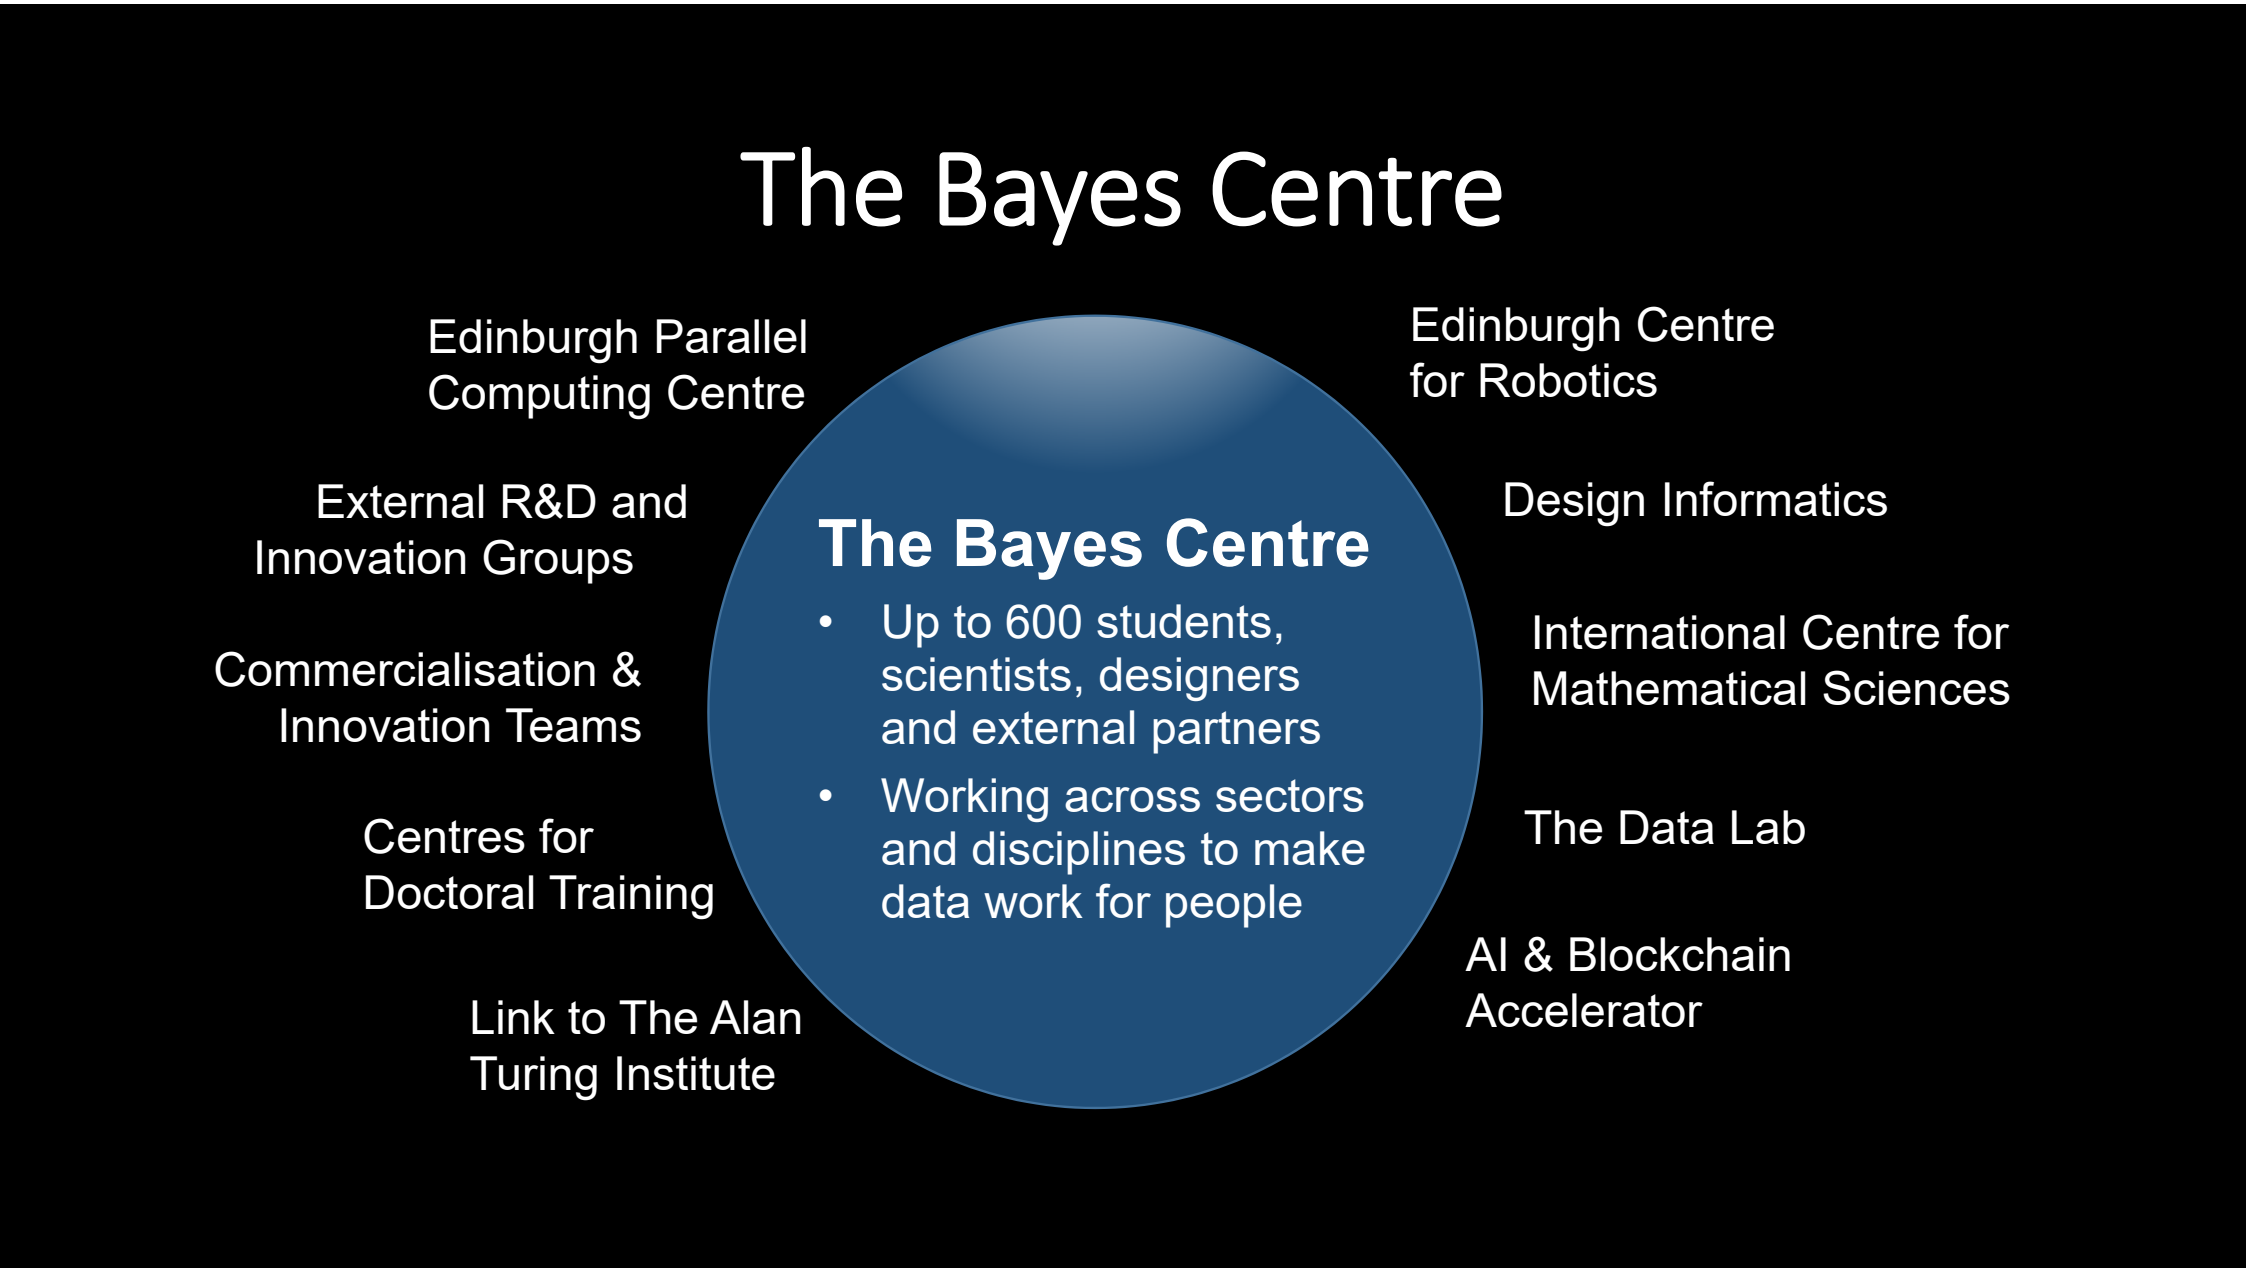

Data-Driven Innovation

City Deal - Five Hubs – Ten Sectors

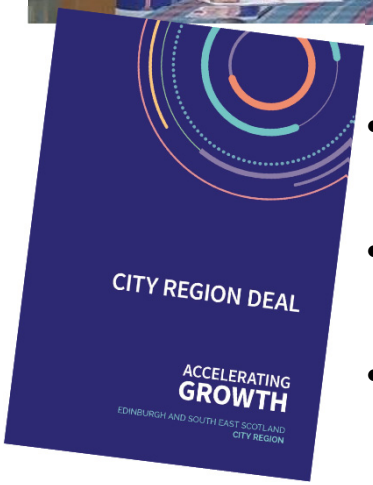

- **2018** – £1.3bn City Region Deal for Edinburgh & SE Scotland
- Includes: housing, transport, skills, culture – and **innovation**
- Data-Driven Innovation component is worth **£660m** and will run for 15 years

Data Science ≠ AI

- Remember the very wide range of areas of AI described earlier. Data driven technologies and probabilistic reasoning are just one part of this range of technologies and applications.

Going Forward – Hybrid AI

Unfortunately there is a history in the subject of the “latest” AI technology predominating and lack of awareness or use of earlier mixed approaches as well as the “hyping” of the latest whiz-bang.

- But for useful “Explainable Systems”…
- Human-machine cooperation & mixed-initiative systems

- Cognitive “human understandable” level
- Sub-cognitive data driven, extensive search space exploration, constraint management level which is explained via a “behavioural envelope” at the cognitive level.

- A good example where a mix of technologies is required to make progress is for…

ROBOTS

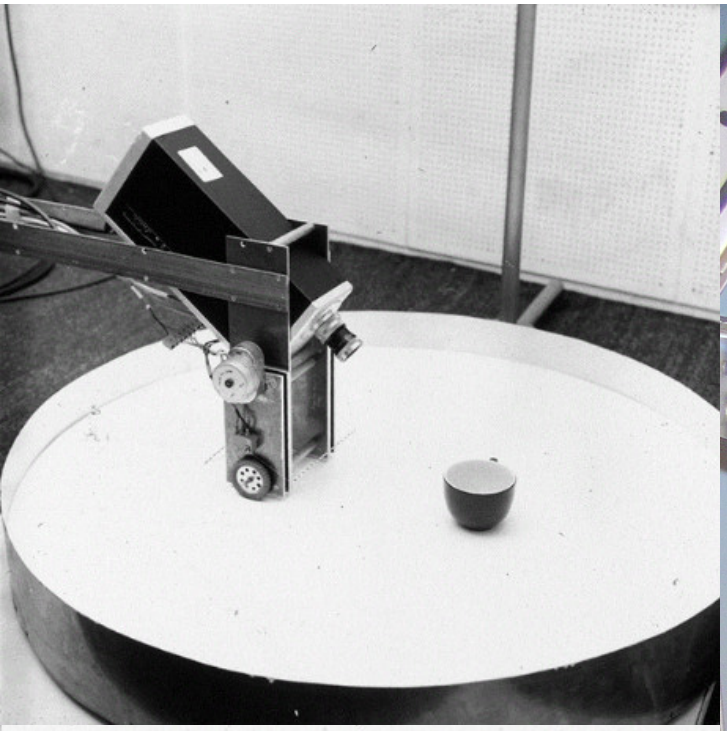

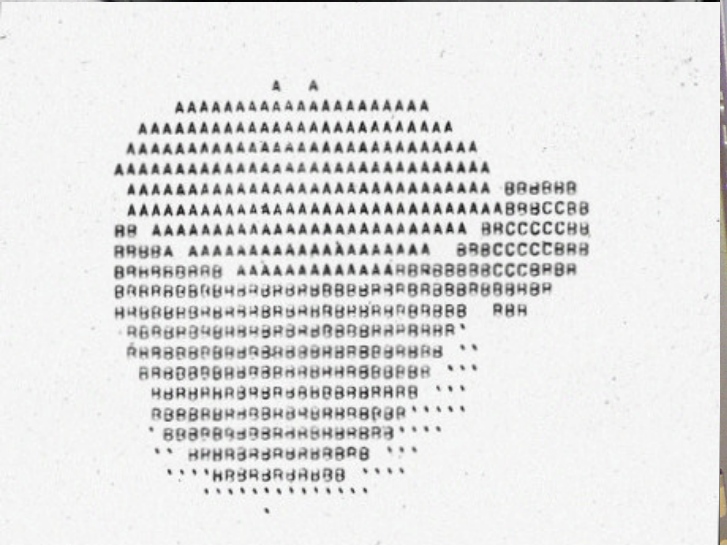

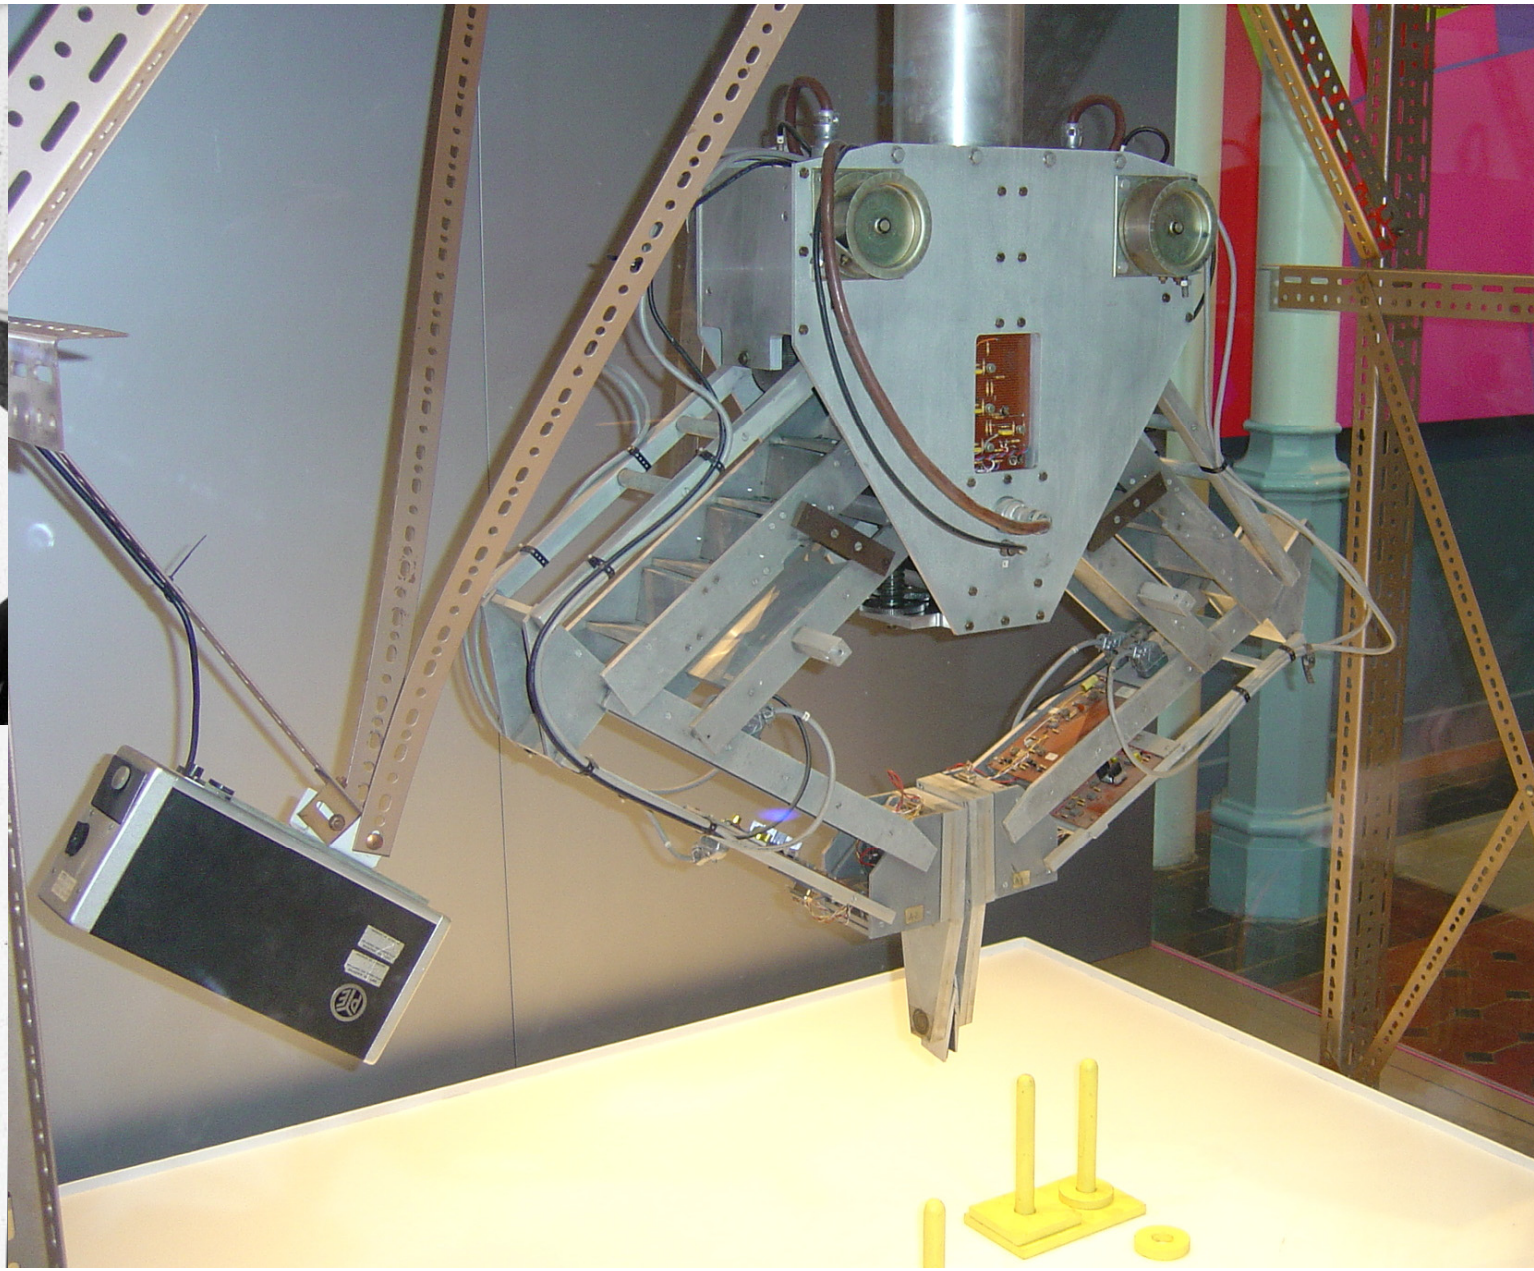

ROBOTS

≠ ≠

Autonomous ≠ Intelligent

Robots ≠ AI

Lower level "reactive" behaviours – becoming very effective due to rapid advances in machine learning and large scale data technologies.

Higher level cognitive capabilities – such as sense making, planning and decision making... often in cooperative situations with humans.

Unimate Robot, 1961


Roomba, iRobot, from 2002

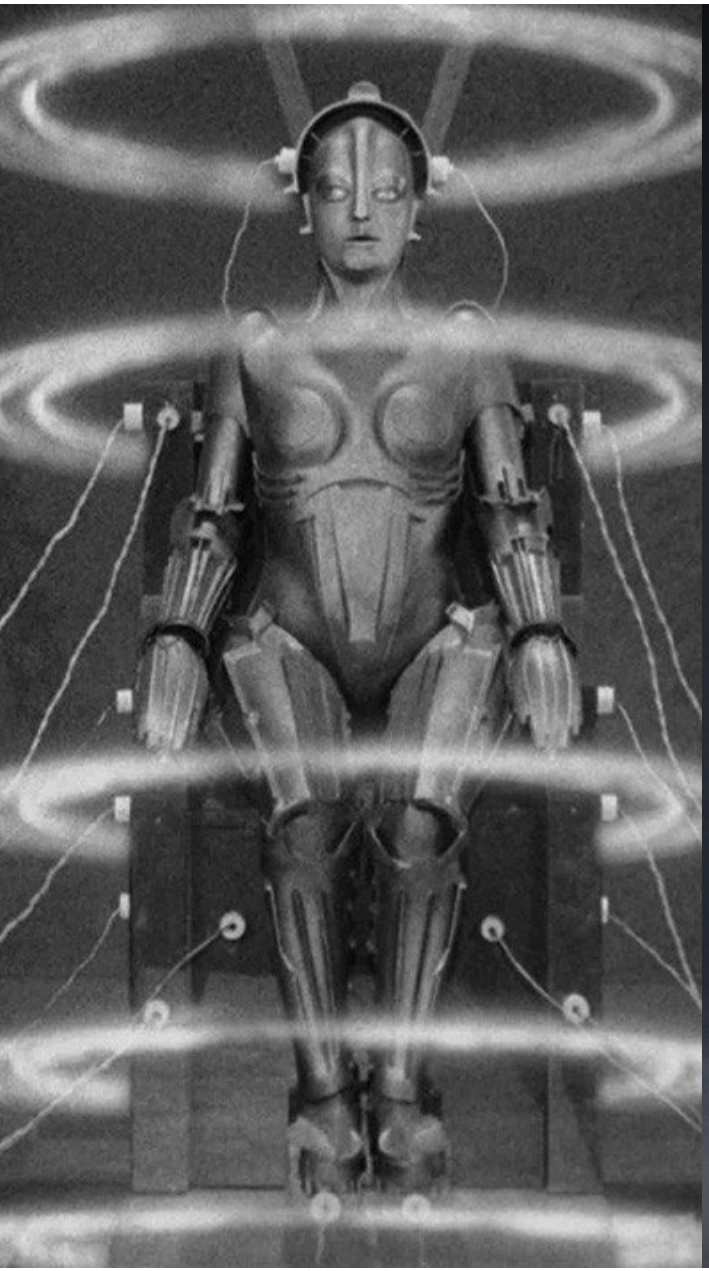

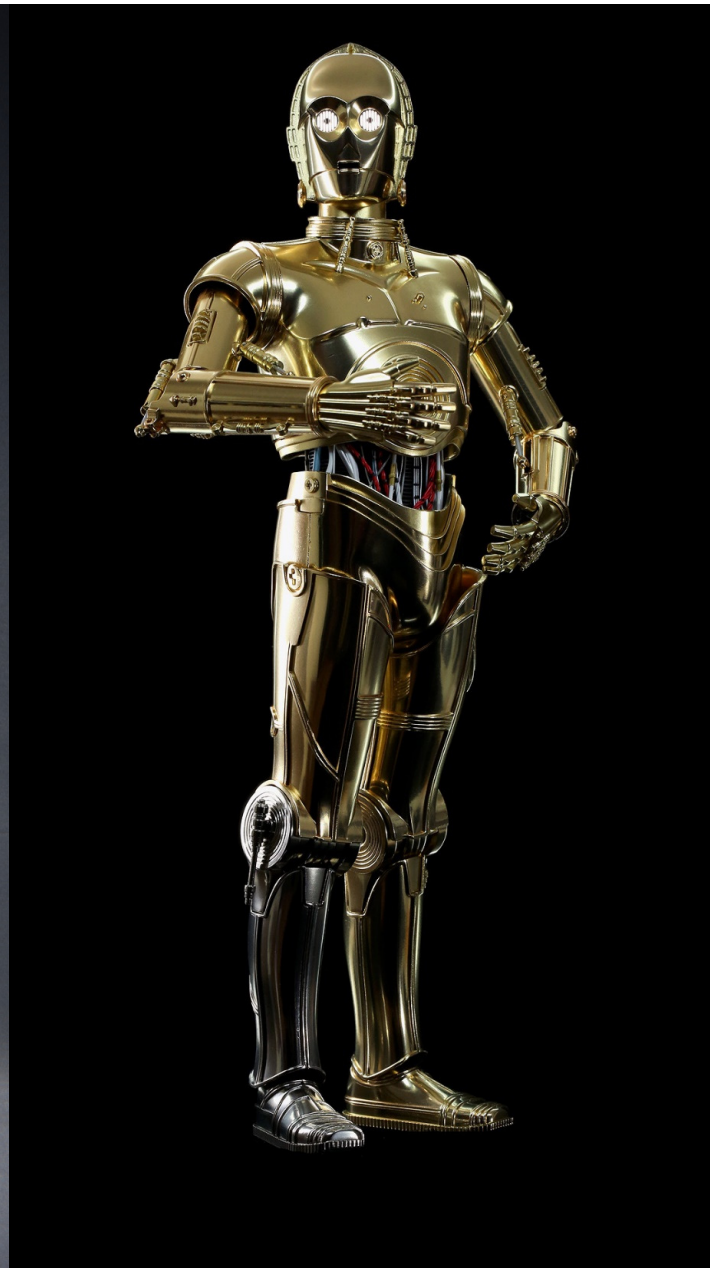

ASIMO
http://www.honda.co.jp/ASIMO/

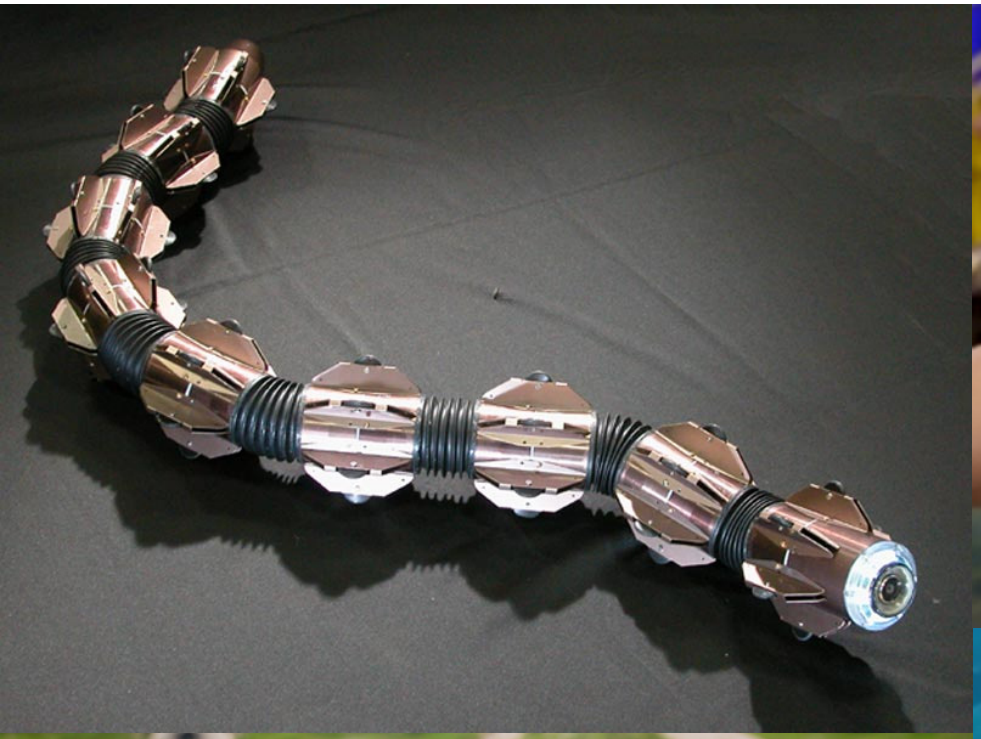

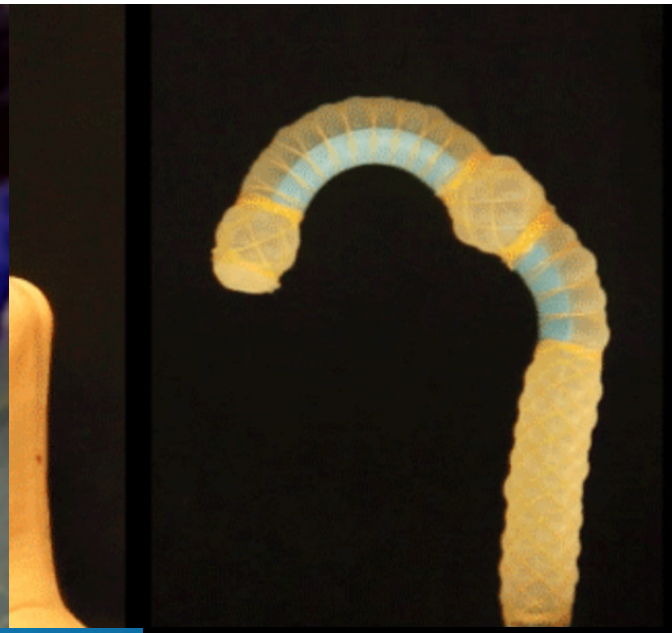


HUSKY A200™
CLEARPATH

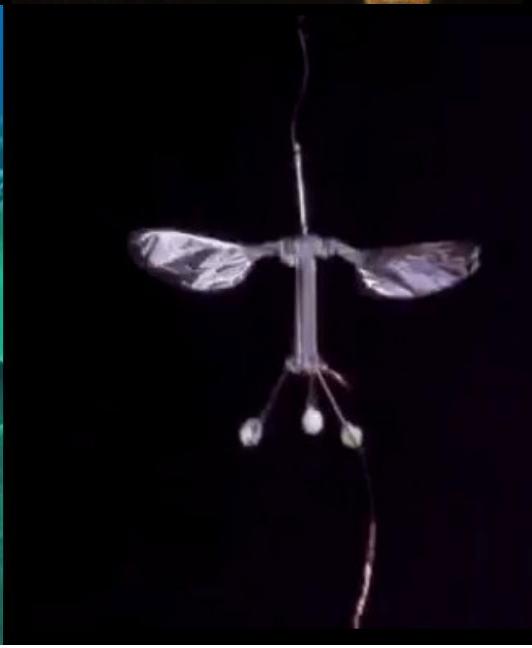

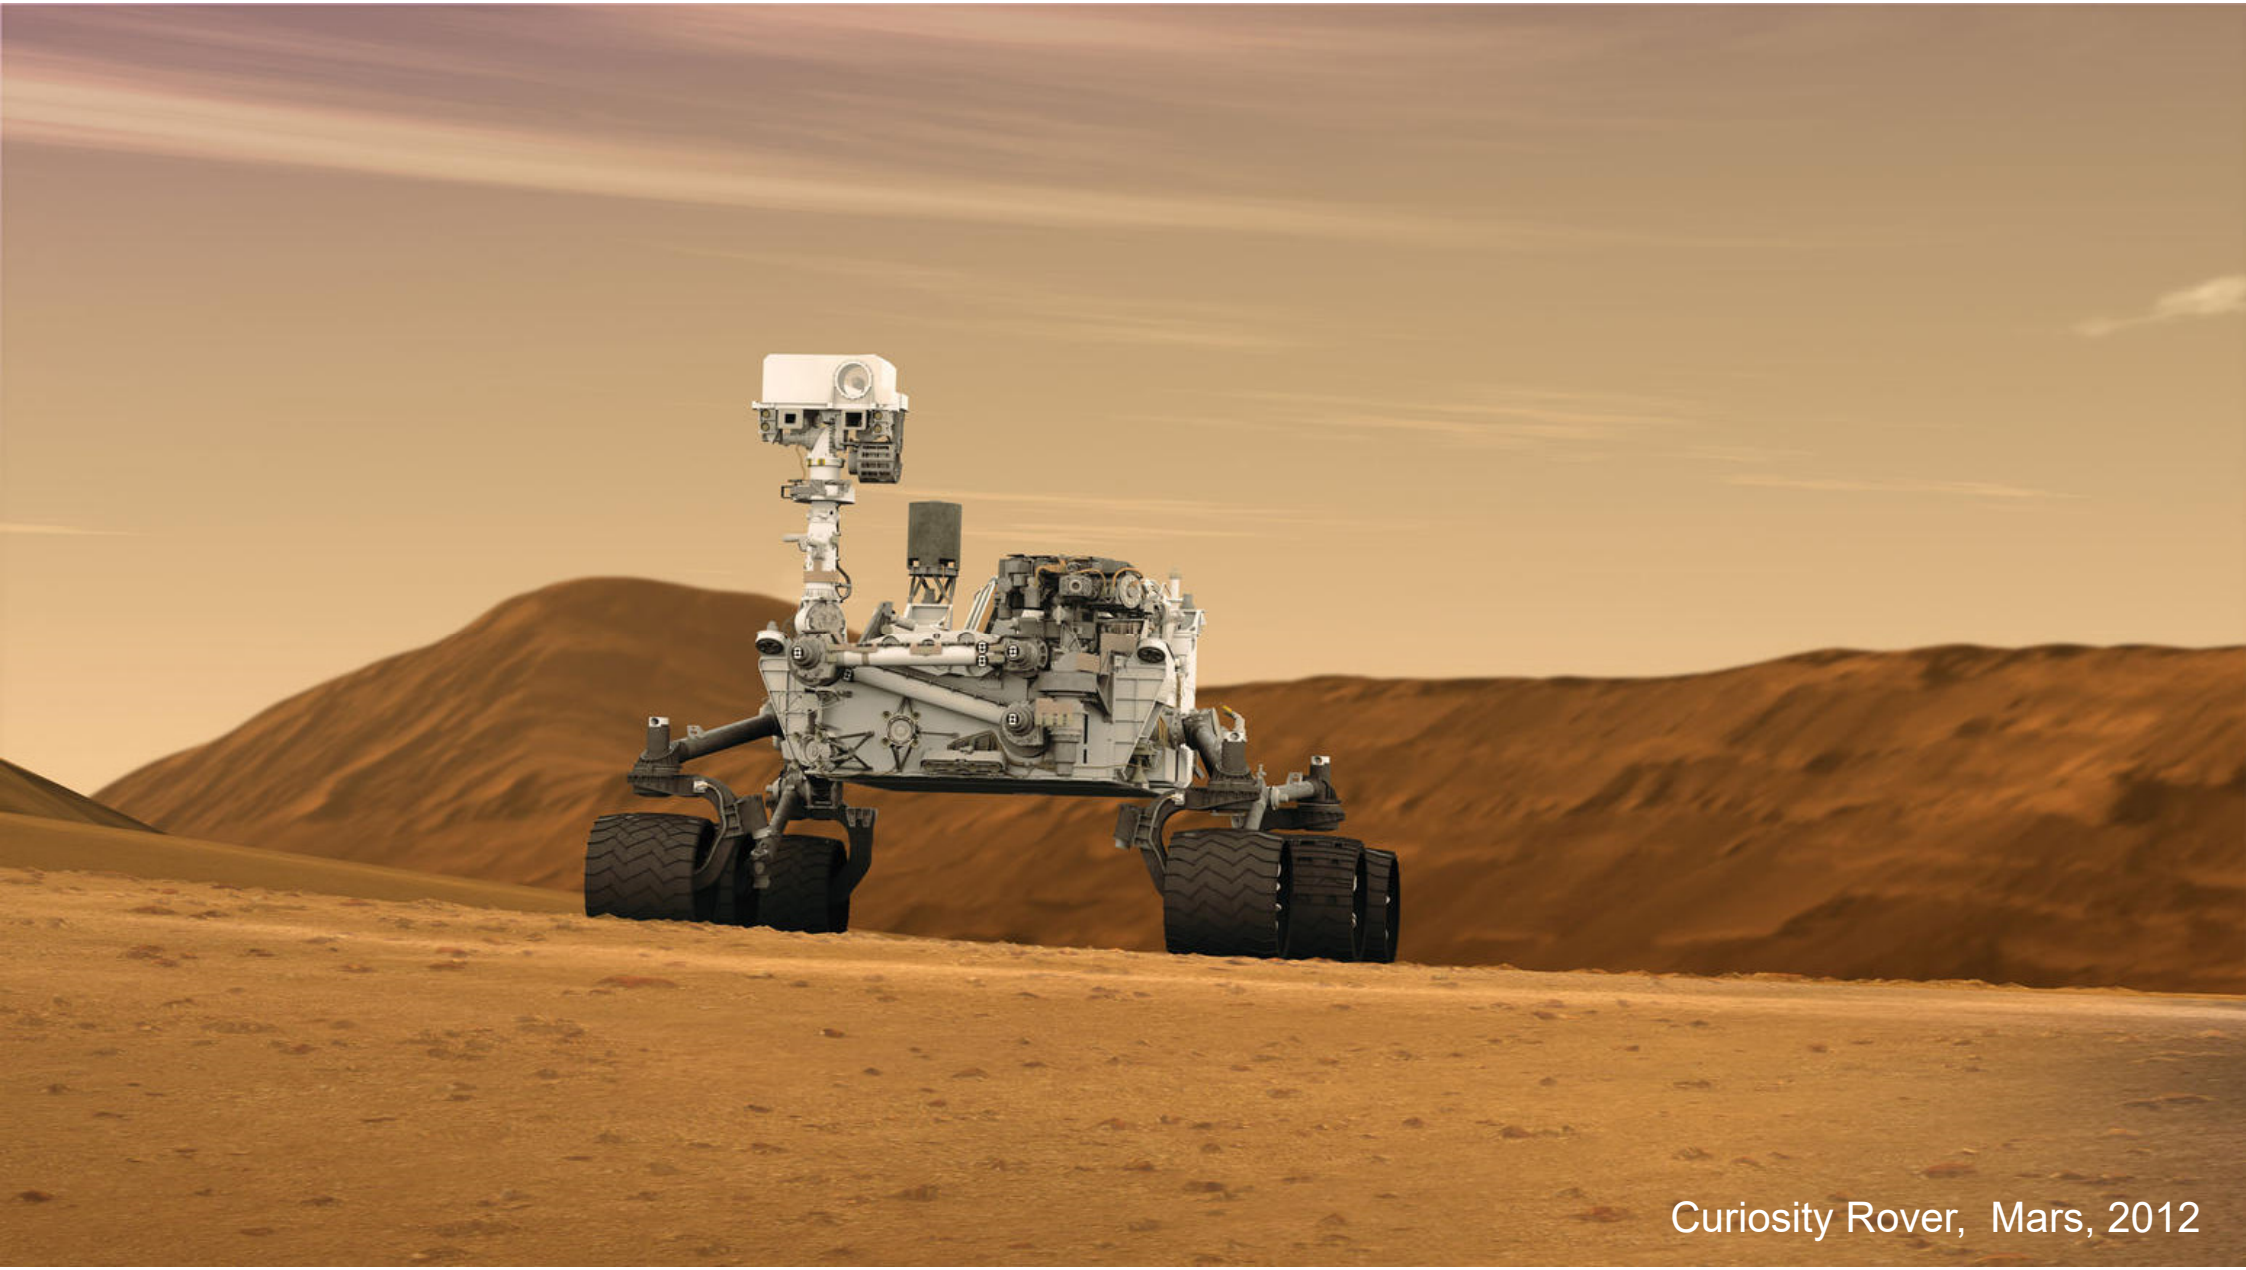
Curiosity Rover, Mars, 2012

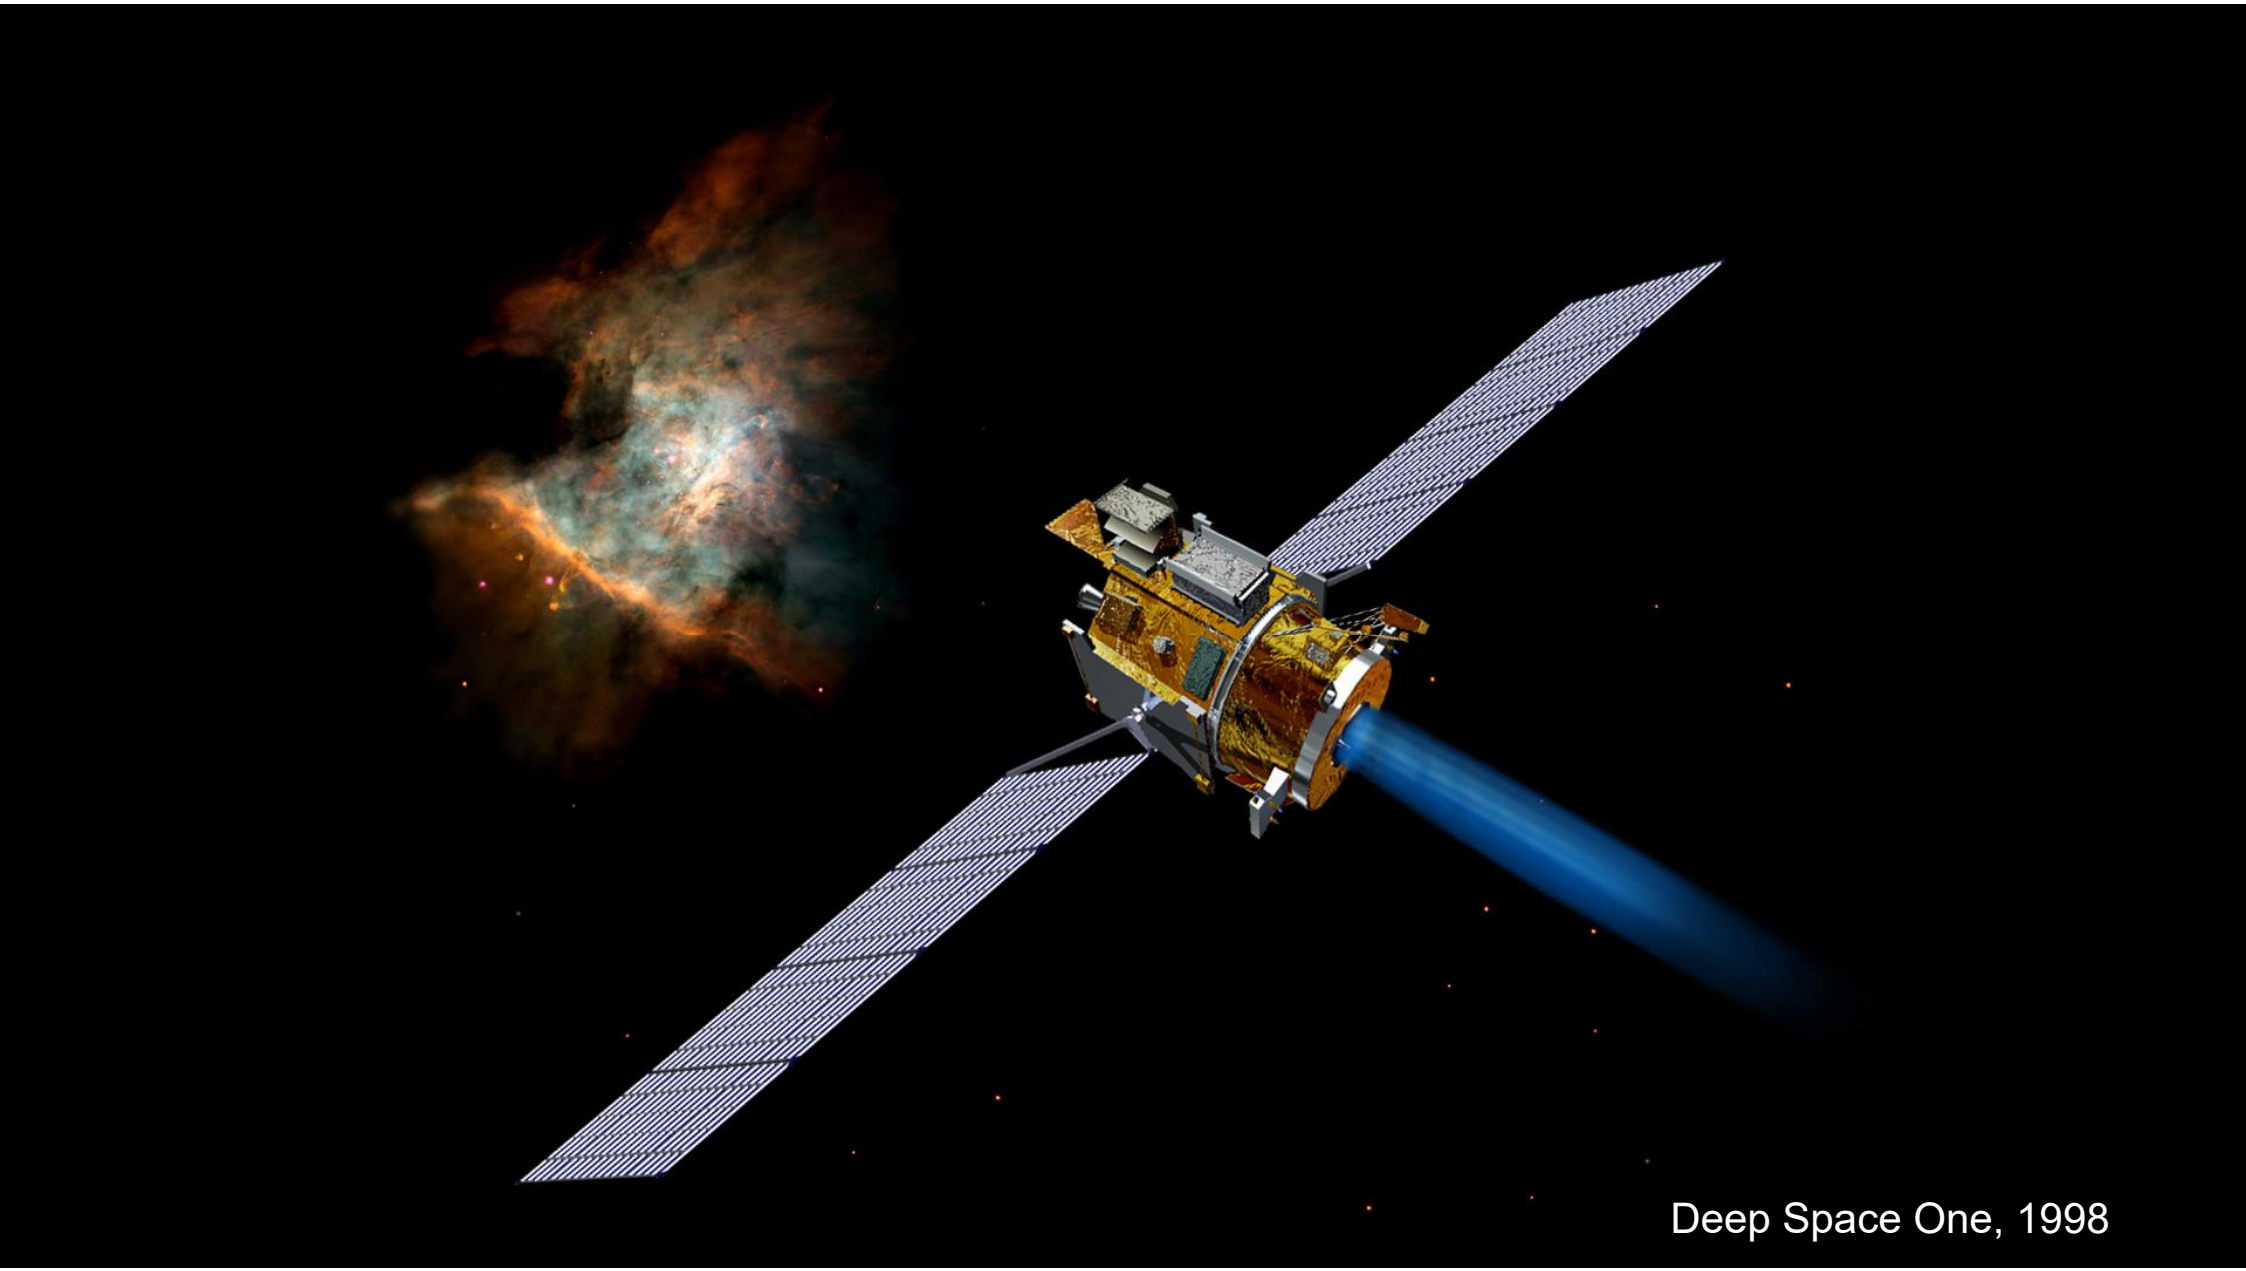
Deep Space One, 1998

da Vinci Surgical Robot, 2000

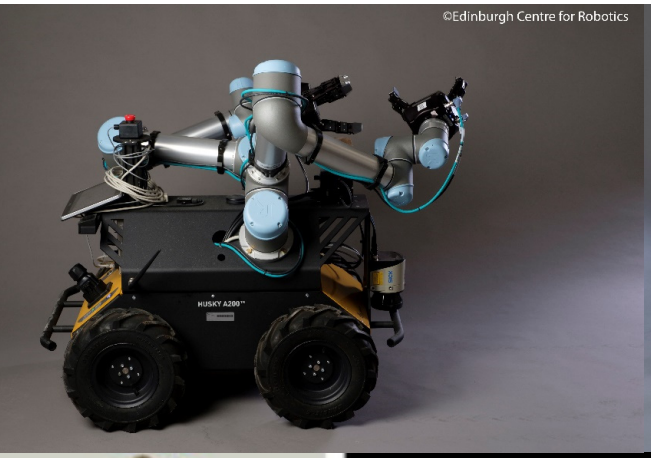

EDINBURGH CENTRE FOR ROBOTICS

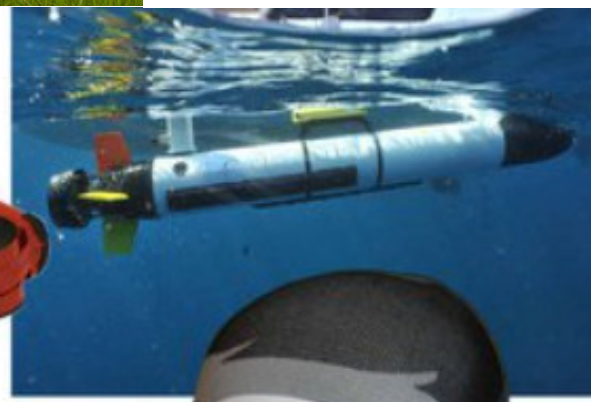

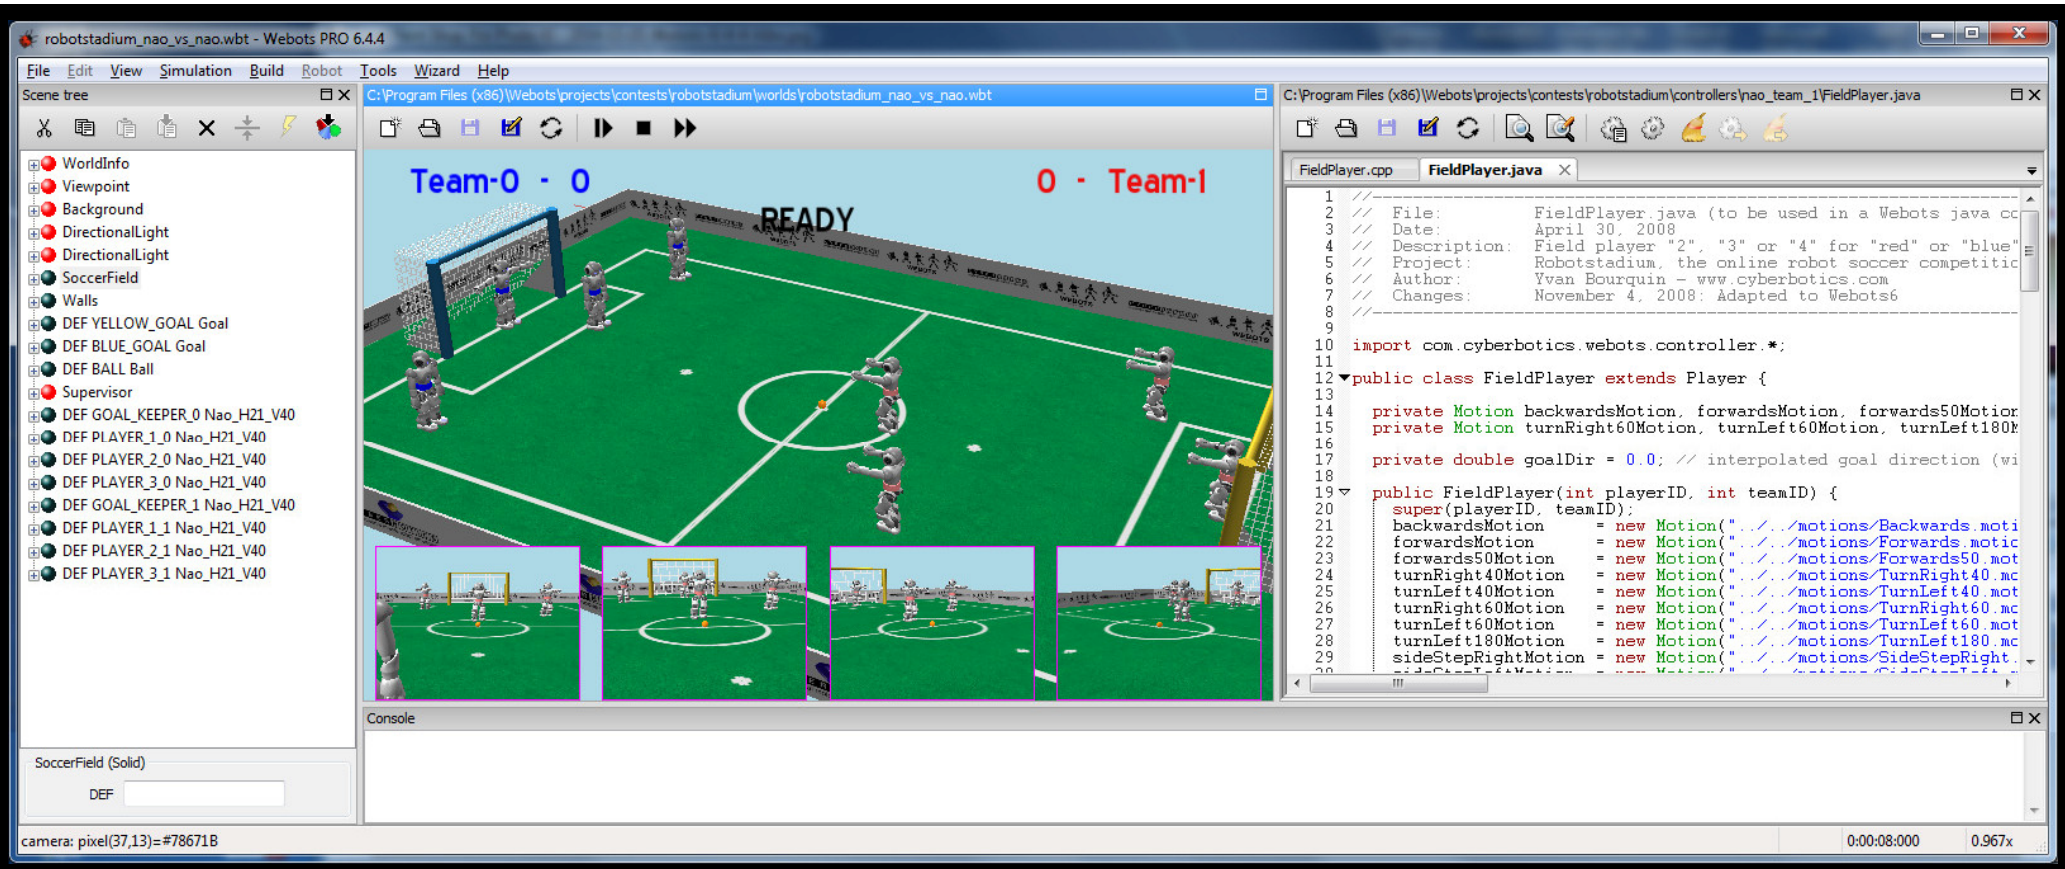
robotstadium_nao_vs_nao.wbt - Webots PRO 6.4.4
File Edit View Simulation Build Robot Tools Wizard Help
Scene tree
WorldInfo
Viewpoint
Background
DirectionalLight
DirectionalLight
SoccerField
Walls
DEF YELLOW_GOAL Goal
DEF BLUE_GOAL Goal
DEF BALL Ball
Supervisor
DEF GOAL_KEEPER_0 Nao_H21_V40
DEF PLAYER_1_0 Nao_H21_V40
DEF PLAYER_2_0 Nao_H21_V40
DEF PLAYER_3_0 Nao_H21_V40
DEF GOAL_KEEPER_1 Nao_H21_V40
DEF PLAYER_1_1 Nao_H21_V40
DEF PLAYER_2_1 Nao_H21_V40
DEF PLAYER_3_1 Nao_H21_V40
SoccerField (Solid)
DEF
C:\Program Files (x86)\Webots\projects\contests\robotstadium\worlds\robotstadium_nao_vs_nao.wbt
Team-0 - 0
READY
0 - Team-1
C:\Program Files (x86)\Webots\projects\contests\robotstadium\controllers\nao_team_1\FieldPlayer.java
FieldPlayer.cpp
FieldPlayer.java
1 //---
2 // File: FieldPlayer.java (to be used in a Webots java cc
3 // Date: April 30, 2008
4 // Description: Field player "2", "3" or "4" for "red" or "blue"
5 // Project: Robotstadium, the online robot soccer competitic
6 // Author: Yvan Bourquin - www.cyberbotics.com
7 // Changes: November 4, 2008: Adapted to Webots6
8 //---
9
10 import com.cyberbotics.webots.controller.*;
11
12 public class FieldPlayer extends Player {
13
14 private Motion backwardsMotion, forwardsMotion, forwards50Motion
15 private Motion turnRight60Motion, turnLeft60Motion, turnLeft180M
16
17 private double goalDir = 0.0; // interpolated goal direction (wi
18
19 public FieldPlayer(int playerID, int teamID) {
20 super(playerID, teamID);
21 backwardsMotion = new Motion("../../motions/Backwards.moti
22 forwardsMotion = new Motion("../../motions/Forwards.motic
23 forwards50Motion = new Motion("../../motions/Forwards50.mot
24 turnRight40Motion = new Motion("../../motions/TurnRight40.mc
25 turnLeft40Motion = new Motion("../../motions/TurnLeft40.mot
26 turnRight60Motion = new Motion("../../motions/TurnRight60.mc
27 turnLeft60Motion = new Motion("../../motions/TurnLeft60.mot
28 turnLeft180Motion = new Motion("../../motions/TurnLeft180.mc
29 sideStepRightMotion = new Motion("../../motions/SideStepRight.
Console
camera: pixel(37,13)=#78671B
0:00:08:000
0.967x

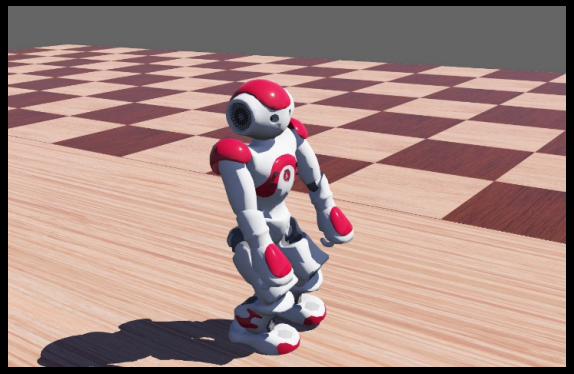

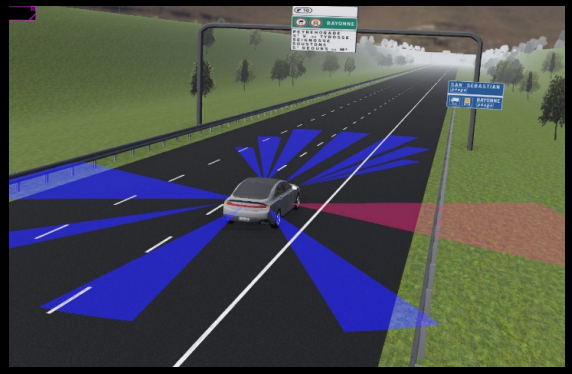

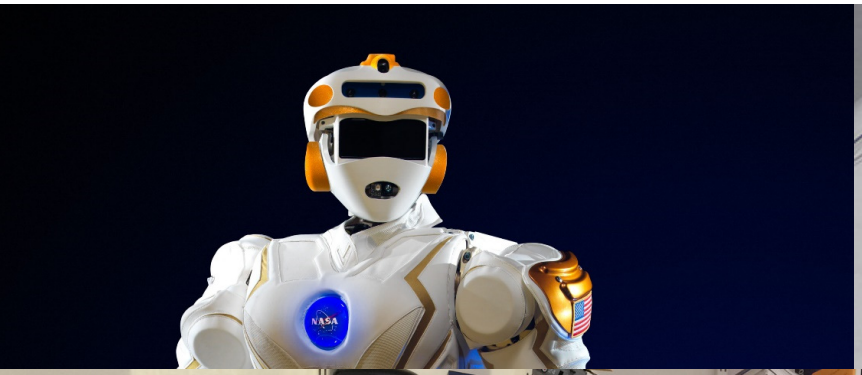



Oculus
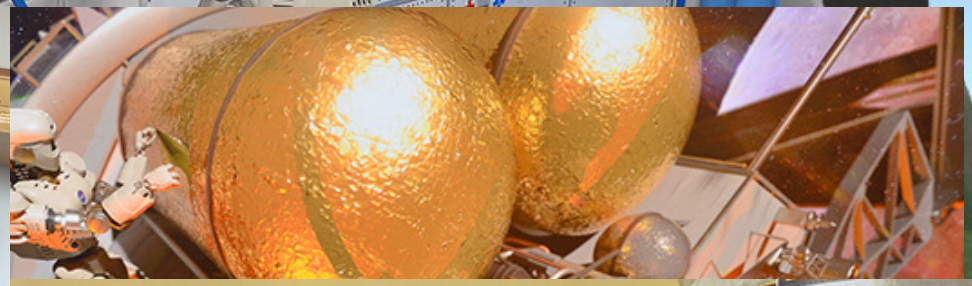
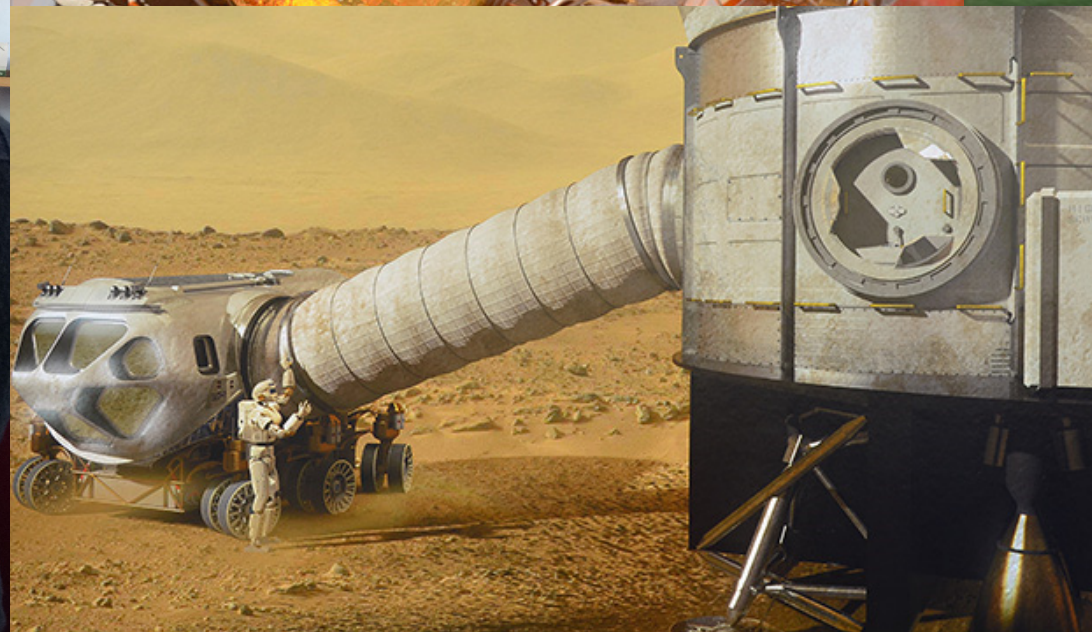

Ethical and Societal Issues

- Privacy issues
- Personal data (mis-)use
- Ease of reversing pseudo-anonymization
- Cultural, racial and gender issues in training data sets
- Autonomous vehicle deployments
- Autonomous weapons deployments - who is responsible?
- Job displacement and work practices
- Wealth sharing and fair distribution


HM Government
National AI Strategy
ALGORITHM
NATIONAL AI STRATEGY – AI SKILLS & TALENT – NHS AI LAB – APPLIED AI
Ethics and safety in AI
A remarkable time of human promise has been ushered in by the convergence of the ever-expanding availability of big data, the soaring speed and stretch of cloud computing platforms, and the advancement of increasingly sophisticated machine learning algorithms.
Innovations in AI are already leaving a mark on government, by improving the provision of essential social goods and services from healthcare, education, and transportation to food supply, energy, and environmental management. These bounties are likely just the start.
The prospect that progress in AI will help government to confront some of its most urgent challenges is exciting, but legitimate worries abound. As with any new and rapidly evolving technology, a steep learning curve means that mistakes and miscalculations will be made and that both unanticipated and harmful impacts will occur.
In order to manage these impacts responsibly and to direct the development of AI systems toward optimal public benefit, The Alan Turing Institute's public policy programme partnered with the Office for Artificial Intelligence and the Government Digital Service to produce guidance on the responsible design and implementation of AI systems in the public sector.
The guide, Understanding Artificial Intelligence Ethics and Safety, is the most comprehensive guidance on the topic of AI ethics and safety in the public sector to date. It identifies the potential harms caused by AI systems and proposes concrete, operationalisable measures to counteract them. The guide stresses that public sector organisations can anticipate and prevent these potential harms by stewarding a culture of responsible innovation and by putting in place governance processes that support the design and implementation of ethical, fair, and safe AI systems.
The guidance is relevant to everyone involved in the design, production, and deployment of a public sector AI project: from data scientists and data engineers to domain experts, delivery managers and departmental leads. Our aim – and hope – in writing the guide is to encourage civil servants interested in conducting AI projects to make considerations of AI ethics and safety a first priority
David Leslie
Ethics Theme Lead, and Ethics Fellow, The Alan Turing Institute
52
Biased & Discriminatory
Unjustifiable outcomes
Unsafe, Unreliable
Organisations can anticipate and prevent potential harms by stewarding a culture of responsible innovation and by putting in place governance processes that support the design and implementation of ethical, fair, and safe AI systems.
Disintegration of Social Connections
No Recourse, No Rights
Invasion of Privacy
Read the Understanding artificial intelligence ethics and safety guide at turing.ac.uk
53


http://www.aiai.ed.ac.uk/~ai/resources/