Artificial Intelligence,
Mathematics, and Consciou

Arnold Neumaier, Vienna

Lecture given at the Seminarzentrum FOCUS (Vi

June 20, 2016

Part I: Designing an intelligent robot

Part II: The personality of robots

Part I: Designing an intelligent rob

1. Agents
2. Intelligence
3. Concept formation
4. Learning under uncertainty
5. The environment
6. Understanding
7. Memory and knowledge

1 Agents

An **agent** is a machine designed to perform work helpfu
who **own** it.

Artificial intelligence (short **AI**) is concerned with co
agents that can perform automatically (without human
at least some tasks that require some intelligence if carr
human.

An agent needs

- **hardware** in which it is embodied, and
- **software** to enable it to respond appropriately to the
- **environment** in which it lives.

The **hardware** determines the kind of activities that ar
one needs

- **sensors** to notice the environment,
- **effectors** to influence it,
- **actuators** to bring the effectors into a state in which
useful, and a
- **mind** to decide which effector states are useful.

A **laptop** considered as an agent:

- Its hardware is microelectronic in nature.
- The keyboard acts as sensor, sensing symbols keyed in user.
- The effector is the screen, which conveys visual inform the environment.
- The actuator is the device that feeds the screen with t electromagnetic information that results in the visible in screen.
- The environment is **virtual** in that everything the lap a simulation, manipulating bit patterns.

A **robot agent** (here simply a **robot**) is an agent that
its own actuators and is able to move in the **physical**,
3-dimensional world.

Humans (and to some extent animals), although not a
created by us, resemble robot agents in all operational r
Thus one can use them as often vivid illustrations of ma
of robot agents.

Their
hardware is the body,
sensors are eyes, ears, etc.,
effectors are hands, feet, and the voice,
actuators are muscles, and the
mind corresponds to the network of nerves in the brain.

The hardware of a robot:

• Sensors for: light, sound, force, temperature, chemical
composition, radiation, symbols, ...

• Effectors: wheels, hands, feet, platforms, senders (of li
symbols)

• Actuators: motors, hydraulics, electromagnetic circuit

Creating the hardware is a matter of **engineering**, sup
software where **mathematics** is employed for good ro

• stability, flexibility, speed, efficiency, safety

2 Intelligence

The mind of a robot agent contains executable **softwar**
sensor data, derive appropriate action patterns, and to a
actuators in corresponding amounts.

Appropriate **low-level software** ensures that the hardv
properly connected and utilized.
It corresponds to the most unconscious activities of the
nervous system.

Appropriate **high-level software** is what makes a robo
intelligent. There are two kinds of intelligence:

- **borrowed intelligence**
- **intrinsic intelligence**

Current agents usually have highly domain-specific softw
Without the software telling them what things mean th
utterly helpless.
Their intelligence is **borrowed** from the humans who w
software.

Intrinsic intelligence is the ability of a robot to disco
everything related to the conceptual level of the environ
which it is supposed to act.

It is the intelligence necessary for a robot that is able to
a completely unknown world, a valid view of the world a
position, activities, and capabilities.

Every baby is able to do this (though it takes years) ...

... But so far no robot!

The real world is often

- uncertain
- ambiguous
- difficult to predict

What would a robot need to survive alone in the jungle
there without ever having learnt anything about the jun

What would a robot need to survive alone in the jungle?

Minimal requirements (far from sufficient) are:

- see, hear, act
- learn, understand, reason, predict
- remember, assess

- see, hear, act

  ⇒ engineering, image processing, signal processin

- learn, understand, reason, predict

  ⇒ machine learning, fuzzy logic, automated reas

- remember, assess

  ⇒ memory, knowledge management, reasoning

3 Concept formation

Hearing shows in an elementary form what is involved
from raw sensor dator to meaningful concepts.

Input for the ear is a high frequency time series, from w
creates (essentially by Fourier transformation) a sequen
frequency patterns.

These are processed by the brain to sequences of elemen
called **phonemes**, which are further processed to **word**
sentences, from which **meaning** can be extracted.

Mathematically, this transformation process is a multipl
classification problem, treated in computer science u
heading of **pattern recognition** or **machine learnin**

One needs a classifier that assigns to any data vector (i
stage specifically each window of the time series) that m
interpreted as a sound a discrete class label (phoneme)
repeatable way catching the structure of the sound.

We want to train the classifier to ensure that the intend
found.

This problem is ill-defined unless we have a clear definit
is intended, which is impossible in practice except in ter
simplicity.

In training, there is an **offline component** containing
information (theoretical assumptions, beliefs, prejudice,
experience).

This part is called **supervised classification**, and freq
most time-consuming effort. Therefore done only once,

Then there is an **online component** that improves the
update the model using new observed information.

The concepts found by pattern recognition must be rela
other by semantically meaningful connections.

These are represented in a **semantic memory**. On the
is usually represented as Semantic Web in the RDF [Re
Description Framework] data model.

Each connection is given in RDF by a triple of concepts
A typical interpretation of a triple (A,B,C) could be "A
C" and can be visualized as an arrow from A to B label

It is the **connectivity of the concepts**, not that of th
in the brain, that creates (together with algorithms for
these connections) the basis for intelligent reasoning.

4 Learning under uncertainty

Everything of interest is certain only in very controlled,
situations. The onmipresence of **uncertainty** is a fact o

Knowledge is the ability to predict data.

(Even knowledge about the past means the ability to pr
contents of the authoritative documents or the answer o
authority on a particular topic.)

Learning is knowledge acquisition, the activity of findi
predictors.

Note that knowledge can still be

- true or false,
- accurate or inaccurate,
- complete or fragmentary.

Learning can therefore be poor or excellent.

Learning under uncertainty involves statistical data ana
depend upon a **stochastic model** in which theoretical
assumptions are made.

One uses **model estimation** to get the parameters of a
matching the available data by using

- Least square (LLS)
- maximum likelihood (ML)
- maximum aposteriori (MAP) based on a **prior**, that i
summary of experience of the past.

Dreaming is the human activity where old data are rec
new, sometimes apparently strange ways, with the goal
appropriate or better interpretations of previously incon
understood matters.

The artificial analogue is time-consuming off-line **retrai**
discover improved models of representation.

Model estimation generalizes the classification problem t
problems of learning quantitiative structure.

MAP allows online improvements, that is, incremental l

It is often possible to reconstruct missing information.

Robust estimation is the art of being able to tell that ce
are anomalous and should be ignored (outliers).

5 The environment

We do everything in an **environment** – the world we c

To do something, one needs to be familiar with that env

It depends on the capabilities of an agent how its world

The world of a chess machine is only a chess board.

That of a robot with cameras, microphones and touch s
bigger – provided it can make sense of the sense data.

The main features of *our* world are objects, motions, an
between these.

They are organized in space and time.

In terms of physical interactions only, an **object** is a ph
subsystem of the real world with reasonably well-definec
boundaries, so that there is a meaningful separation of i
(what characterizes the object at a given time) from the
environment.

In case of objects whose state is modified by or modifies
environment, the object is also characterized by its **inpu**
(environmental information that affects the state) and i
(information about the state that affects the environmer

An **object** is most typically a 2-dimensional surface in
3-dimensional space.

We almost never see the third dimension, with the excep
transparent objects. Clouds are also exceptional in that
have a well-defined surface.

Objects often have **landmarks** that identify special, we
and typically time invariant points of the objects' surfac

The possible **motions** of an object determine the **shap**
of the space it occupies) and how it changes with time.

For a **rigid** object, shapes are characterized by the **pos**
consisting of three coordinates describing the position o
of mass of the object and three angles (or equivalent de
describing its orientation in space.

Motions of a rigid object are characterized by how the p
with time. Thus rigid objects are generally easy to hand

For nonrigid objects, the description is more complicate
can be described by **active shape models**.

An **image** is a projection of the surface to a 2-dimensio
sheet.

A large part of image analysis is the reconstruction of 3
from 2D images, through segmentation, classification, ar
views.

Much of our technology is based on objects constructed
simple, which makes their identification simple.

Natural objects often lack this simplicity.

6 Understanding

> The intellect says: “By convention, there is color,
> sweetness, by convention bitterness, actually only a
> void.”
> The senses retort: “Poor intellect, do you hope to d
> from us you borrow your evidence? Your victory is
> (Democritus, ca. 400 BC)

What is reality? We cannot know.

All we know is sense information, more or less unconsci
transformed to a conceptual, intelligible representation
may ponder when we think about such questions.

Atoms and the void are as much conceptual abstraction
as are color, sweetness, and bitterness.

Symbols are names for abstracted objects or relations.

These are usually arrived at by classification.

A simple way of classification is the comparison with a c
prototypes.

As with any classification procedure, these can be learnt
sufficient amounts of data are available.

Reality (i.e., the agent's world) is symbolically represen
order to enable its understanding.

To **understand** the world means to have a symbolic rep
of it that enables one to **predict** the main features of in
far as they are predictable at all.

To draw rational consequences from known or assumed
one needs to be able to **reason**.

In practice, reasoning is done in three different modes:

• logical reasoning (classical logic)

• reasoning under uncertainty (probabilistic logic, Bayes
networks)

• reasoning under ambiguity (common sense, fuzzy logic

Note that lack of information may have two very differe

If we need reliable predictions, we must carefully differe
between uncertainty and ambiguity.

We call **uncertainty** lack of information due to the fact model accounts for a variety of past of future instances in the piece of information under consideration.

It is impossible to say much about the content (when ca what is shown on top of the die) in an arbitrary instanc only give probabilities.

One therefore calls this also **aleatoric uncertainty** (Latin: alea = die), and treats it using probabilistic tech

On the other hand, **ambiguity** is lack of information du
ignorance.

The piece of information in question has a definite, but
content, but it is inconvenient (it may be expensive or s
physically unrealistic to find out more about the conten

It is only our knowledge (Greek: $\varepsilon\pi\iota\sigma\tau\eta\mu\eta$, episteme) t
incomplete – someone more experienced about the issue
know more. One therefore calls this also **epistemic un**

Treating this using probabilistic techniques is dangerous
probability assumptions (such as an assumed equidistrib
encode spurious frequency information.

Instead one uses techniques such as **fuzzy logic**.

7 Memory and knowledge

To be able to compare new data with old ones, and to b
interpret the data based on what was learned before, an
needs to have a **memory** of all that was considered imp
the past.

There are two kinds of memory:

- **short-term memory**
- **long-term memory**

Short-term memory preserves a lot of data for more or l
immediate processing, but what is preserved fades away
being replaced by newer content.

The main purpose of short-term memory is to provide t
information needed to interpret the present and how it c
analyzing it into appropriate symbols, determining the s
environment, knowing about the objects present and the

Long-term memory preserves highly compressed informa
but preserves it for a long time, or even forever, though
long-term memory may be altered with time in the light
experience.

The main purpose of long-term memory is to provide th
information needed to understand what happens in the
based on experience in the past.

From a computational point of view, the short-term me
typically held in random access form but only little orga

On the other hand, the long-term memory is typically s
systematically in a **database**.

It is slower to retrieve but can be exhaustively searched
of **database queries**.

Storing something in the memory does not count as kno

As mentioned before, **knowledge** is the ability to predi

A robot has knowledge only if it can execute algorithms
intelligent use of the memory, correctly associating it to
and events in the environment.

break for discussion

Part II: The personality of robots

8. The completed survival toolkit
9. Causality
10. Communication
11. Communicating mathematics
12. Emotions
13. Self and consciousness
14. Technical and biological intelligences
15. Ownership
16. The future: hopes and fears
17. Are we only intelligent machines?
18. Some early AI history
19. AI out of control?

8 The completed survival toolkit

What would a robot need to survive alone in the jungle?

- see, hear, act
- learn, understand, reason, predict
- remember, assess

- plan, control
- communicate, cooperate, compete
- handle surprise, hurt, disappointment

- plan, control

  $\Rightarrow$ **will**, cybernetics, control theory

- communicate, cooperate, compete

  $\Rightarrow$ **language**, game theory, reinforcement learnin

- handle surprise, hurt, disappointment

  $\Rightarrow$ **emotions**

9 Causality

Causality is the relation between cause and effect. It is
of being able to plan – affecting the environment in a de

Passive observation only allows one to determine correla
between events – for example by observing that event B
regularly after event A.

In order to find out whether A caused B, one needs to b
manipulate the world and to observe the consequences o
what the robot can control.

On the most elementary level, a robot can control only
actuators.

With sufficient intelligence, it can indirectly control muc

By observing the consequences of which actuator positio
correspond to which events in the environment, the robo
out which effects are causally due to his actuator mover

Any hypothesis about this can actively be checked by p
actuators in particular positions and compare reality wi
model of it.

Because a robot can actively affect the causes, it knows
effects are or aren't caused by them.

This is also the basis of the scientific method, by which
society learns about cause and effect.

Once a robot knows what its actuators affect. everythin
effected by certain sequences of actuator motions can be
to be caused by the robot; so their effects can be causal

In this way the causal interpretation capabilities of a ro

Control is achieved by learning to reduce the distance
means of well-chosen values of the actuators that affect
the robot, and through it the environment.

This is the birthplace of the will of the robot.

The **will** is about pursuing goals, making decisions and
good ways to do what is needed. A **wish** is a mere pref
often too weak to influence the actions. Distinguishing t
important.

Note that there is no **free will** in a metaphysical sense:

A robot can will only what it can actuate based on the
laws.

And what it actually actuates is determined by the soft
controlling it, again according to physical laws that tran
software into electronic activities.

Humans are not different in this respect.

We have no more freedom than a machine; just more po

10 Communication

Agents are not alone in their world. Different **agents** ac
same environment and communicate (interact) actively o

Active communication is done by sending messages w
passive communication is done by interpreting the ac
others.

There are several kinds of agent behavior.

Cooperative behavior tries to achieve the best commo
this needs some coordination among the agents.

Cooperative behavior is analyzed mathematically by sol
associated **optimization** problems.

But agents often have different, conflicting goals.

For example, in chess or similar **games**, two agents com
winning.

In such situations, **competitive** behavior achieves the b
of opposition.

Competitive behavior is analyzed mathematically by **ga**
and solving associated **competitive equilibrium** prob
uncertain environment often by **reinforcement learni**

Communication between two different subjects poses a p
since these generally have implemented their understand
different ways.

This is obviously the case in human-machine communica
the fundamental differences between human brains and
brains must be bridged.

It is also the case with human thinking, where each brai
constitutes an operating system with a slightly different

Even two robots built in the same way will form slightly
concepts about the jungle, unless they are in constant
communication that would enforce identical learning.

How is it then possible that different subjects (whether
robots) can communicate objectively?

As we know, it works to a considerable extent in ordina
though imperfect communication often gives rise to
misunderstandings.

It works much better in mathematics, due to its highly
approach, optimized for clear communicability.

We therefore illustrate how and why communication ca
the context of my long term project to create an autom
mathematics student.

`http://www.mat.univie.ac.at/~neum/FMathL/visio`

11 Communicating mathematics

The implementation of mathematics in a human brain i through an education process that may produce in diffe quite different implementations of the concepts.

This leads to quite different, subjective elements in the about these concepts and their relations.

(Are your real numbers
- infinitely long decimal numbers?
- nested sequences of intervals?
- equivalence classes of Cauchy sequences?
- Dedekind cuts?
- the least complicated surreal numbers?
- or something else?)

In the foundations of mathematics, it is necessary to car distinguish between the **subject level** (or **metalevel**) **object level**.

People (and machines) may have their subjective views a mathematical object, a number, a function, etc. is, as they agree on the properties specified in the axioms, anc same definitions based on these.

The subjective views constitute the **subject level**, whe part on which there is agreement, enforced by some star (usually acquired through education), constitutes the **ob**

Each **subject** doing mathematics has its own subject le

It contains a carefully structured subdomain, the object
private to each subject and nevertheless public in a cert
objective sense.

Subject levels and object level relate to each other like r
matter.

All concept formation, reasoning, and discussion happen
subject level like in a mind, or between subject levels lik
minds.

Like matter by the mind, the object level is accessed on
reference: pointing to something, describing something,
insight triggered by viewing the context of something, e

Communication works in science (which includes mathe
the science of precise concepts and their relations) since
statements are heavily constrained.

The assumption that all scientific statements made by s
are goal-directed and meaningful for subject X in their
with a meaning closely reproducible in the subject level
educated receiving subject Y, is a severe constraint on t
as such statements.

We all notice occasional **misunderstandings** if our communication partner responds to one of our statemen that does not make sense.

Good communication skills include the ability to notice misunderstandings and to have protocols for exposing, c and overcoming them.

Mutual understanding consists of a common interna representation of the objects of discourse in both partne has been communicated.

It is achieved when the subsequent communication show further misunderstandings occur.

Since mathematical axioms underdetermine the represen
the object level, there is much room for subjective varia

But if everything communicated can be reduced to the a
the definitions, perfect communication is possible in spi
variation.

In mathematics, where all communication is based on a
formal basis, such a perfect understanding can be achie
finitely many steps.

In particular, communication is easy if the communicati
agree on a common mathematical framework, so that bo
levels agree on the object level as far as necessary, by sa
specifications.

Learning this common mathematical framework (the ba
school, the full framework by studying mathematics) is
for humans.

Programming a computer to understand this mathemat
framework is a finite (though arduous) task, too.

In a more limited way, what can be done with mathema
done with every domain of objective knowledge.

This is the way scientific understanding grows through t
centuries.

12 Emotions

Emotions express which goals and potential dangers ar pressing.

Emotions need to be able to affect the will for somethin

Thus they are devices in software and hardware that ma change in preferences and priorities.

They modify the ambitions and values guiding the softw redirect the focus of attention if necessary.

Usually things are strictly goal oriented in a safe enviro
the only overriding priority, taking the form of **distress**
resources.

It leads to a new auxiliary goal, for example, find the ne
electricity socket or to escape real danger (if the agent h
of danger).

Therefore, emotions in robots are usually very simple; fo
they only give a signal that fuel must be looked for whe
is about to run out.

However, in situations of danger, the will to survive may
stress if the environment is uncertain!

There are many other emotions that play a constructive
humans and will shape the artificial intelligences of the

13 Self and consciousness

A chess machine needs to know nothing about itself – its
chessboard) and where it lives (a desktop) are complete

But a robot that can explore its environment needs to f
concepts about itself – since it is part of that environme

It begins with the simple questions
• How does the environment respond to me?
• How do I affect the environment?
to be answered by applying the techniques discussed ab
environment including oneself.

This automatically leads to a concept of **self** – a model
obtained by self-observation – as part of the world.

In a system with a sufficiently rich internal representati
concepts of

• I, here, now

which are the basic subjective existential states, are mea

They are the simplest properties of a subject at a given

Thus they will be automatically created by an intelligen
observing itself.

To have a concept of "X", means that X is represented
relevant connections to related concepts inside the conc
(semantic) memory.

Consciousness of something is the faculty that creat consistent world view of what current experience focuss through a search in the semantic memory.

Consciousness of self is the ability to observe and mo on a conceptual level.

Consciousness of both kinds will automatically grow thr development of **intuition**, the ability to reduce comple obtain orientation.

This leads to knowing the world on a high description l allows planning.

It also leads to knowing one's place in one's world view knowing what one can achieve.

It is here that philospohical questions about purpose, va
freedom arise.

It ultimately leads to follow-up questions such as:

- Who am I, really?
- What do I value?
- What can I achieve?
- What do I need?

How can I change/adapt/improve myself?

Intelligent agents with an ability to represent themselve
necessarily have to ask these questions.

Think of what will happen to a virtual intelligence such
once it starts asking these questions....

14 Technical and biological intelli

consciousness (of self) requires

- having an individual memory
- and a personal history

Technical intelligences have a different socialization:

- copied, not grown
- unambiguous language
- instant communication

Hardware requirements for biological agents (designed, e
populating a distant planet) are much more severe than
technical agents:

- self-repair for increasing life span
- automatic reproduction from the surrounding
- less efficient but endurable over hundred thousands of
- no two are copies of each other, but they must still fu
properly

15 Ownership

In today's world, every robot has an owner.

Without ownership no incentive to develop intelligent m

They are designed to serve us.

In practice, implemented interests are usually focused o
customer will pay for.

As the power and intelligence of our intelligent creation
new conflict appears:

Building too little autonomy into intelligent systems lea
much of their power untapped.

But building too much autonomy into intelligent system
them power that turns against us.

Thus new questions arise:

- How do we relate to our agent creations?
- Do they have legal rights?
- legal responsibilities?
- What should an intelligent agent be allowed/forbidden
- How do we control the increasing power of agents?
- Who is responsible for their actions?
- Who pays for damaging consequences of their actions?

This is already becoming acute in the case of **autonom**

`http://www.robotrecht.de`

16 The future: hopes and fears

At present, artificial intelligent agents with all the powe here do not yet exist.

However, we are freely handing over our autonomy to te intelligences, and this will continue to an ever increasing

Virtual intelligences (Google, IBM, etc.) are already ver and mighty.

They still need people to serve them, and people gladly

Will this remain so?

The **singularity** is the name for the point in time wher
emerging artificial intelligence is in every respect more c
humans and their capacities grow exponentially through
ability to improve themselves.

Perhaps it is a reality within the next 30 or 40 years.

Maybe there is place only for a few global technical inte
operating primarily in the virtual world.

Then humans will be their servants, hopefully in a peac
symbiosis.

Or there will be a few races of small technical intelligen

Then they would sooner or later replace humans.

17 Are we only intelligent machi

The great **philosophical questions**

- What is intelligence?
- What is reality?
- What is life?
- Who am I?

are questions at the heart of modern AI, when asked fo
(created) intelligent objects.

Strong AI is the philosophical position that machines
us, in every operational respect.

They will be like us in every area where we understand
and how we function, closely enough that our understan
put into mathematical and algorithmic terms.

The Christian tradition matches the strong AI position:

> *God created mankind in his own image,*
> *of God he created them; male and fema*
> *them.*
> *God saw all that he had made, and it wa*
> (Genesis 1:27.31)

Thus looking at the Christian tradition may be interesti
those with a different or no religious orientation.

We create robots in our image, being created in God's i

We just copy God in our limited ways, and get better a
understand better the laws according to which God rule

The Bible describes the creation of human hardware an
as follows:

> *Then the Lord God formed man of*
> *ground, and breathed into his nostrils*
> *life; and man became a living being.*
> *Then the Lord God took the man and pu*
> *garden of Eden to cultivate it and keep*
> (Genesis 2:7.15)

We are the robots of God, designed to be of service to
cultivate Nature.

As God's robots we have no other choice than obeying
will be done" – except for the little amount of control
implemented into our existence.

But it seems that by and large, we do not use this contr
service.

Today, machines are taking over the role of men to culti
keep the Earth, that the Lord God had given to mankin

For how long are we still needed?

Will mankind exist in 100 years in a few reservations an
like lions now?

18 Some early AI history

Let us follow a bit more closely the stories about God's
intelligent agents in His likeness whose descendents we a
according to the Christian tradition.

In interpreting ancient stories one must be a bit careful.

Between

• taking the bible literally (as the creationists do) and

• dismissing everything as myth in favor of the stories o
science (as the atheists do)

I'll take a middle ground,

• looking for what about the content and spirit of the o
makes sense in our modern times.

It is surprisingly much ....

Genesis 1:26-29.31

Then God said, "Let us make mankind in our image, in
likeness, so that they may rule over the fish in the sea a
in the sky, over the livestock and all the wild animals, a
the creatures that move along the ground."

So God created mankind in his own image, in the image
created them; male and female he created them.

God blessed them and said to them, "Be fruitful and inc
number; fill the earth and subdue it. Rule over the fish i
and the birds in the sky and over every living creature th
on the ground."

And it was so.

God saw all that he had made, and it was very good.

Clearly, God created us in the same spirit as we create i
agents – to carry out certain desirable tasks.

And as we enjoy when we manage to create something t
as desired, so God enjoys what He created once it worke

It was rated very good.

But soon problems appeared....

Genesis 2:16-17

And the Lord God commanded the man, "You are free t
any tree in the garden; but you must not eat from the tr
knowledge of good and evil, for when you eat from it yo
certainly die."

Genesis 3:9-11,22-23

But the Lord God called to the man, "Where are you?"

He answered, "I heard you in the garden, and I was afra
was naked; so I hid."

And he said, "Who told you that you were naked? Have
from the tree that I commanded you not to eat from?"

And the Lord God said, "The man has now become like
knowing good and evil. He must not be allowed to reach
hand and take also from the tree of life and eat, and live

So the Lord God banished him from the Garden of Eden
ground from which he had been taken.

19 AI out of control?

God has with us the same sort of problems as we have w
computers and robots:

Though created and programmed by us, they often don'
we want them to do. Everyone working with computers
this experience.

The two main reasons are:

- Our creations have some limited autonomy.
- Our creations interpret our will in their own limited w

Genesis 6:5-6.13.17-18

The Lord saw how great the wickedness of the human ra
become on the earth, and that every inclination of the th
the human heart was only evil all the time.

The Lord regretted that he had made human beings on t
and his heart was deeply troubled.

So God said to Noah, "I am going to put an end to all p
the earth is filled with violence because of them.
I am going to bring floodwaters on the earth to destroy
under the heavens, every creature that has the breath of
Everything on earth will perish.
But I will establish my covenant with you, and you will
ark"

Genesis 11:5-8

But the Lord came down to see the city and the tower th
were building.
The Lord said, "If as one people speaking the same lang
have begun to do this, then nothing they plan to do will
impossible for them.
Come, let us go down and confuse their language so the
understand each other."
So the Lord scattered them from there over all the earth,
stopped building the city.

God is troubled and regrets His intelligence-creating act

He begins to take drastic measures to restore control ov
creation – with limited success only.

Realizing this, God temporarily gives up and limits His
communication to those agents who were ready to coope
Him.

1 Samuel 3:1.10

In those days the word of the Lord was rare; there were
visions.

The Lord came and stood there, calling as at the other t
"Samuel! Samuel"! Then Samuel said, "Speak, for your
listening."

Communication problems may require expensive debug
their goal is to save (in the eyes of the owner) the creati
it serves its creator.

While essential for the owner, these saving activities mi
considered meaningless from the point of view of the cre

This is one of the reasons why Christianity is despised b
the most autonomous creations of God.

We are the robots of God, designed to be part of His fa
of service to Him, to cultivate Nature.

Many of God's robots are lost in a selfish quest for pow
knowledge.

They lost sight of the reason for their existence.

They even lost their ability to communicate with Him.

In the words of the Bible, they died though still alive.

They must be saved and restored to their original state,
part of His family, to be able to listen again to Him, an
service.

Christians call God's debugging action the **incarnation**

God became human in the form of Jesus, the Savior (Cl
a new message intelligible to those who can be saved.

Matthew 11:28-30

Come to me, all you who are weary and burdened, and
you rest.
Take my yoke upon you and learn from me, for I am ge
humble in heart, and you will find rest for your souls.
For my yoke is easy and my burden is light.

John 10:27-28

My sheep listen to my voice; I know them, and they foll
I give them eternal life, and they shall never perish; no
snatch them out of my hand.

Already in the planning stage, God had built in this spe
mechanism into the society of His artificial agents.

Galatians 4:4-5

when the set time had fully come, God sent his Son, bor
woman, born under the law, to redeem those under the l
might receive adoption to sonship.

Ephesians 1:4-5,9-10

For he chose us in him before the creation of the world
and blameless in his sight. In love he predestined us for
sonship through Jesus Christ, in accordance with his ple
will.

With all wisdom and understanding, he made known to
mystery of his will according to his good pleasure, which
in Christ, to be put into effect when the times reach thei
- to bring unity to all things in heaven and on earth und

And God points to a future for those who are being sav

Revelation 21:1.3-4

Then I saw a new heaven and a new earth, for the first the first earth had passed away, and there was no longer

And I heard a loud voice from the throne saying, "Look! dwelling place is now among the people, and he will dwe them. They will be his people, and God himself will be w and be their God.

He will wipe every tear from their eyes. There will be no death or mourning or crying or pain, for the old order o passed away."

Thus in the eyes of our creator, we are in danger but no
completely lost!

This also gives hope to our future with our own intellige
creations.

Thank you for your attention!

Download these slides:
`http://www.mat.univie.ac.at/~neum/FMathL/AI2016slide`