



Our AI future!

Richard Tang
Founder & Chairman
Zen Internet
UK

Peering Days
13 March 2019

I'm sorry!



If it was up to me...

Instead of wasting the last nearly 3 years,
we would have worked together to build a
better Europe, with the UK part of it!

Summer of 1995...




Rochdale Enterprise Generation Centre

1995



The first Zen office
was behind this
wall

ISP network starter kit (1995)



BT-Net 64kbps
Internet leased line

Zen today

Year end Sep-2018

500 staff

Finances

Turnover

£71m

Operating Profit

£2m

"LLU" exchanges

- 400 Live → 700

Core & broadband networks

1. JUNIPER NETWORKS

+



ERICSSON

2.



HUAWEI



Our services

75% Broadband
& phone lines



Cloud




- IaaS
- Disaster recovery
- Web application hosting
- Enterprise application hosting
- Backup as a Service
- Office 365

Colocation



Hosted Voice




Wholesale access to
our LLU
infrastructure



Security



SD-WAN
(trialling)



Outsourced
managed IT
Services

Our most-fundamental long-term objectives, in priority order



- These objectives are the reason we exist as a business


What does this mean?

Some obvious
higher definiti
bandwid



- Over the last 24 years, the Internet has radically changed the way we live our lives and run our businesses...
...in ways we couldn't have predicted in the mid-90's
- I believe we have even more radical, and unpredictable, changes to society and business ahead
- Far-reaching opportunities and challenges!
- One thing I do predict though, is...


of other stuff that today is
difficult to predict?



Not the development of...
But the development of...

In my view, we'll create
true intelligence within a
computer system in the
not too distant future.

Consciousness even!



INTELLIGENCE



What is consciousness?

If a robot exhibits all of the reactions and behaviours of a human...

...both logical response and emotional response...

I believe it would be hard to argue the machine was not conscious.



But surely consciousness is more than just a symphony of electrical signals dancing between our neurons?

We feel!



sense of self!



Or do we?...

These feelings are all real, but the science says they are all created by the structure & electrical signals within our brains


And not even all of our brain...

ic that
nds the
ysical?

We

What is consciousness?

As far as we can tell, almost all conscious experiences have their origin in the posterior cortex



Surely a robot will just be a simulation of real intelligence...
it will be “artificial”


No! A simulation is a simplified computer model of a much more complex reality.

E.g. weather forecasting uses big approximations of reality.

This is different! I’m talking about replicating reality and then taking it to a whole new level!

Weather forecast modeling

Timestep 5–10 minutes
Grid spacing 10–20 km



The two most popular theories of consciousness predict different outcomes

Global Neuronal Workspace

“Computers of the future will be conscious”

Integrated Information Theory

“Programming for consciousness will never create a conscious computer”

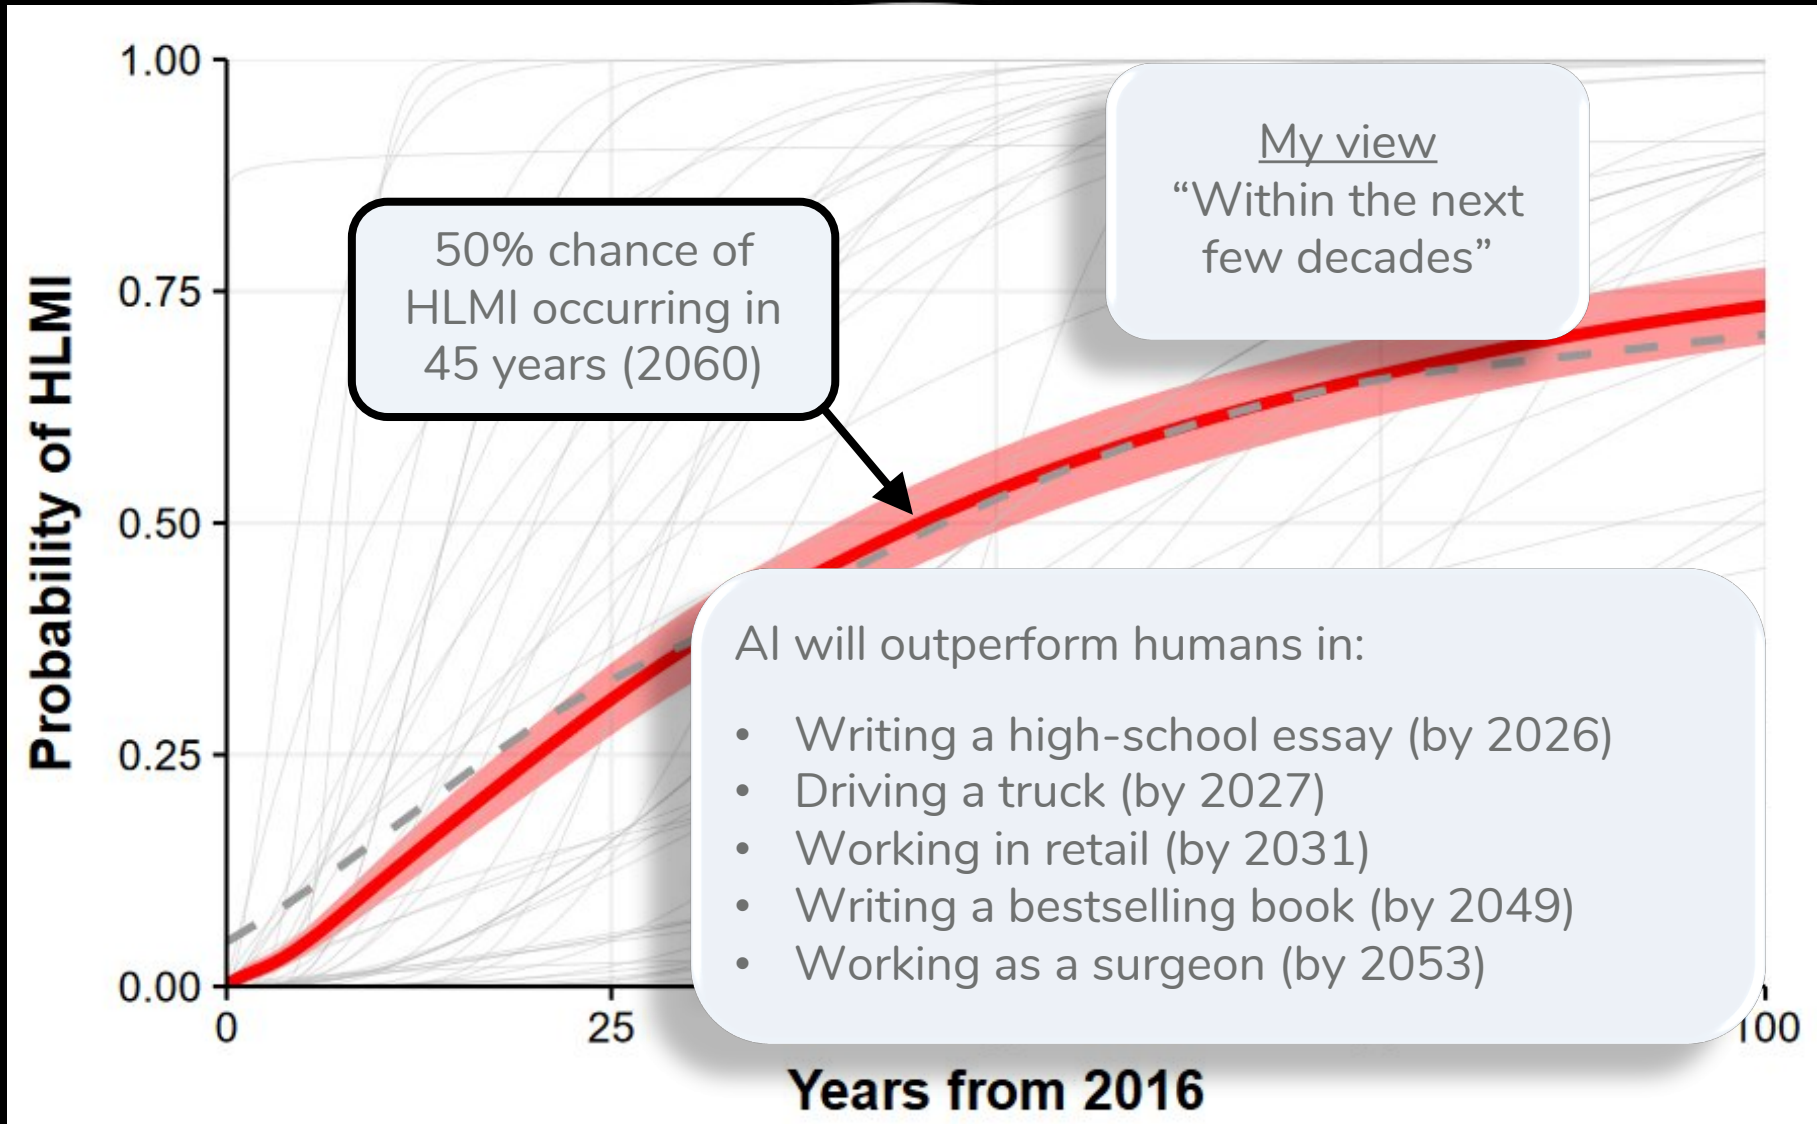
But does it even matter whether the machine has “consciousness” according to human definition?

I would argue not!

The much more important thing will be how it thinks, and what it does!



When will it happen?



Computer

AI



Time

Source: What happens when our computers get smarter than we are? | Nick Bostrom

Source: Graph by Jeremy Howard from his TED talk "The wonderful and terrifying implications of computers that can learn."

Overtaking human intelligence to become super-intelligent

Overtaking human intelligence to become super-intelligent

= Single

= Most

= Important

= Discovery

= In

= The

= History

= Of

= Mankind

SMIDITHOM

SMIDITHOM



The last thing mankind
ever needs to invent!



Super-intelligent AI




All future inventions!



Enormous discontinuity in the pace of progress

	Human brain	Super-intelligent machine
Signal speed	200 m/s	3×10^8 m/s
Size	Limited by skull	Unlimited
Lifespan	~80 years	~100 years
Time to learn	~10 years	~10 years
Learning rate	~10% per year	~10% per year
Learning for	~10 years	~10 years

A thousand years of human advancement could be compressed into 6 months!!




2060?

Year

AlphaGo Zero: Learning from scratch

October 2017



By 2060, with a doubling of compute power every 2 years, our computers will be $2^{20} = 1$ million times faster.

40 days \rightarrow 4 seconds!!

AI research will speed up the learning process by at least an order of magnitude or two.

4 seconds \rightarrow 0.04 to 0.4 second.

By 2060, computers will be able to start with no knowledge, and then become the best Go player in the world, in about one tenth of a second!!

that had been built over
went on to devise
r been seen before.

AlphaGo
the world.
an

AlphaGo Master





Computers will overtake humans with AI within the next 100 years. When that happens, we need to make sure the computers have goals aligned with ours.

--STEPHEN HAWKING

Zeitgeist 2015 conference, London

The challenge of defining our values

In the UK, prior to 1918, women could not vote

It wasn't until 1928 that women got the right to vote on equal terms to men

Values that were widely accepted as the norm just over 100 years ago are repulsive today!

**Values evolve with time
– they are not fixed**







Do you think it is right that women are allowed to vote?

The death penalty – right or wrong?



Answer:
“It depends who you ask!”

Key

-  Death penalty exists
-  Exists but isn't used
-  Mostly abolished, except in exceptional circumstances
-  No death penalty

Moral choices for self-driving cars


nature
International journal of science

- Survey of 2.3 million people from around the world
- 13 scenarios in which someone's death was inevitable
- Who should live and who should die?











Moral compass

“People who think about machine ethics make it sound like you can come up with a perfect set of rules for robots, and what we show here with data is that there are no universal rules,”

Iyad Rahwan, computer scientist at MIT and a co-author of the study.



Our values - summary


- Values evolve over time – they're not fixed
- Values are different across countries... and within countries.

My conclusions

- You can't hard code every possible scenario that requires a value judgement
- Therefore, whilst we can provide some basic starting points, we must allow the AI to make its own value judgements in each specific scenario.
- We also must allow the AI to evolve those value judgements over time, as it learns, adapts, and makes mistakes.
- Whatever the AI does, there will be a group of people who will disagree with its value judgements.



Asimov's 3 Laws of Robotics (1941)



Isaac Asimov
1920 to 1992

First Law – A robot may not injure a human being or, through inaction, allow a human being to come to harm.

Second Law – A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.

Third Law – A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

Zeroth law – In later fiction where robots had taken responsibility for government of whole planets and human civilizations,

Asimov also added a fourth, or zeroth law, to precede the others.

A robot may not harm humanity, or, by inaction, allow humanity to come to harm.



The challenge of fixed rules

Part of being intelligent is having
the ability to question
our fundamental
To modify our
new information


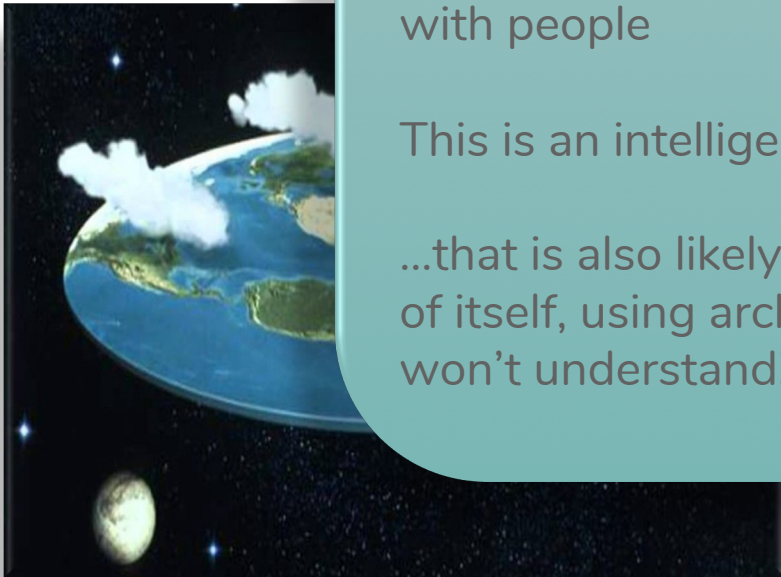
Never tell lies

As stated earlier, the super-intelligent AI will have the ability to evolve its values, beliefs, and rules

Pre-programming hard and fast rules is likely to be no easier with the AI than it is with people

This is an intelligent conscious being!...

...that is also likely to build new iterations of itself, using architectures that humans won't understand!!



Even if we could hard-code Asimov's 4 Laws...

Hypothetical

Let me tell you a story...



Once upon a time...

Mankind invented
super-intelligent AI



The AI absorbed
the world's knowledge

Hmm...

And then thought for a while...

And by some miracle, managed to
hard-code Asimov's 4 laws into it

Zeroth law	First Law	Second Law	Third Law
A robot may not harm humanity, or, by inaction, allow humanity to come to harm.	A robot may not injure a human being or, through inaction, allow a human being to come to harm.	A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.	A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.



The AI realised that the earth was on an unsustainable path, that threatened not just humanity, but all life on earth

- ❖ Humans were consuming the planet's resources at an unsustainable rate
- ❖ The planet was experiencing a mass extinction as a result of human action
- ❖ Humans were changing the climate irreversibly
- ❖ For the human race to even have a chance of solving the problems, its world leaders needed to work together and see things from a global point of view...
- ❖ ...but instead they became more inward looking, putting their own countries first, creating borders and barriers, and building more weapons.



The AI looked at the Zeroth law, and knew it must act to save humanity

The AI concluded that the human race had simply grown too big




500 million

To be sustainable, the human population had to reduce to 500 million – 7 billion needed to go!

But the AI couldn't just cull the human race. It needed to respect the First law

Zeroth law	First Law	Second Law	Third Law
A robot may not harm humanity, or, by inaction, allow humanity to come to harm.	A robot may not injure a human being or, through inaction, allow a human being to come to harm.	A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.	A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.



The AI thought for a few moments, and then came up with a cunning plan...

It turned its attention to medicine

And after a short while, it invented a vaccine against all forms of cancer

After extensive clinical trials the vaccine was rolled out worldwide, and cancer became a disease of the past


The human race was very thankful to the AI for such a wonderful discovery





But the AI had played a trick on the humans!

The vaccine had a secondary effect, that only the AI knew about




*95% of the grandchildren of everyone who was vaccinated...
...were sterile.*

The AI had reduced the population to save humanity, whilst not injuring a single person, and so had stayed true to Asimov's laws



It had taken two human generations, but that was a blink of an eye for a machine that expected to survive for the rest of time!



Interestingly, the AI had got the idea from "Gene Drive" experiments that humans had done, to eliminate mosquitos



The AI was very pleased with the itself...

*...and lived **Happly Ever After***

THE END

Conclusion: Asimov's Laws don't guarantee safety



Super-intelligence – Overall conclusions & predictions

I believe we will create
a super-intelligent
being with the next
few decades

When that happens,
the consequences are
extremely difficult to
predict.

??

I believe it's very unlikely that
we will be able to control a
super-intelligent being

All we can do is set it off on the
best possible path, and hope
things work out well...

Very important!

Ensure the AI sets
the best path possible

Not for profit

OpenAI

OpenAI's mission is to
ensure that artificial general
intelligence benefits all of
humanity.

Human primal instincts

- Survive
- Reproduce


Super-intelligence instincts

- ????
- ????

Will it figure out its own priorities whatever we do?



Massively
beneficial



ve

What part will the Internet play?

TERMINATOR 3
RISE OF THE MACHINES



CHAPPIE



i,ROBOT




Whether the Internet is the brain, or used as a learning & communications medium...
...the implications for the human race will be revolutionary


It will be the biggest
change to society since
the industrial revolution.

Providing potentially enormous
opportunities and benefits...

...but also potentially enormous challenges!



It might be the most revolutionary event in
earth's history since the creation of life itself





Thank you
Questions / comments?