Alexander Kruel

Thoughts and news on transhumanism, vegetarianism, science fiction, science, philosophy, math, programming, language, consciousness and the nature of reality.

‹ What (not) to worry about? • Richard's paradox ›

Randal Koene On Cryonics

2013-01-16 in SI/LW | 1 comment

Related LessWrong discussion: [See here (http://lesswrong.com/lw/fz9/more_cryonics_probability_estimates/82sf)]

Relevant third-party quotes:

> If you don't sign up your kids for cryonics then you are a lousy parent.

— Eliezer Yudkowsky, Normal Cryonics (http://lesswrong.com/lw/1mc/normal_cryonics/)

> I've signed up for cryonics (with Alcor) because I believe that if civilization doesn't collapse then within the next 100 years there will likely be an intelligence trillions upon trillions of times smarter than anyone alive today.

— James Miller (http://lesswrong.com/r/discussion/lw/8f4/neil_degrasse_tyson_on_cryonics/6l2v)

About Randal Koene: [Homepage (http://www.randalkoene.com/)] [FAQ (http://www.randalkoene.com/personal-faq)] [Wikipedia (http://en.wikipedia.org/wiki/Randal_A._Koene)]

Randal Koene On Cryonics

Question: Given your professional expertise in computational neuroscience, neural engineering and information theory I would love to hear your general opinion about cryonics and how likely you deem the possibility that brains frozen given current technology might experience a fatal "loss" of information.

Randal Koene: I guess I should address this by taking the problem apart a bit. We need to determine what we consider "fatal", we need to consider what sort of preservation it would take to recover using technologies we can currently conceive and where we have enough understanding to make credible statements, and then we need to address this other matter of recovery "regardless of losses" due to the supposed extreme superintelligence's capabilities. Finally, I think it is also worthwhile to (at least theoretically) consider what it means to personally deem recovery successful.

1.) "Fatal" information loss:

The gold standard for recovery would probably be a reconstitution that results in no more changes of brain-state than we are ordinarily used to experiencing from moment to moment. Obviously, we are not fixed entities, our brains are plastic, we are always learning and forgetting and changing. Some rate of change seems "normal" to us.
The ideal is a transition in which long-term memory, as well as the intermediate chemical storage through LTP and LTD (long term potentiation and depression), diffuse modulatory states, working and short-term memory instantiated purely through nested patterns of activity in rhythmic oscillations are all captured and then reactivated.
That is not even a goal of current cryopreservation. It is generally understood that the procedure would conserve only long-term states and perhaps LTP&LTD.
How bad is that? Probably not terribly, given that we seem satisfied that we "survive" when recovering from unconsciousness.

A "fatal" loss could depend very much on the area of the brain that is affected by preservation problems. Let's say that the ultrastructure of connectivity in neocortex is somewhat distorted, resulting in a 10% difference between pre- and post-cryo brain configuration. That might not be so terrible…. it is reminiscent of early Alzheimers and in this case one may even improve! But how about functions that are not supported by as much distributed storage as we suppose the neocortex provides? Take dentate gyrus and its hypothesized conceptual pointer function… connectivity changes there could make episodes of memory entirely unretrievable. But still, is that "fatal"?

Personally, I think what I would consider fatal is a change that altered my personality to the point where I would not recognize myself to my own satisfaction (which is of course difficult to test… see point 4). So, a lot of this will be about the circuitry that determines our typical reaction/response, our emotional make-up, our connection to drives and interests. I imagine that a bad cryo-preservation could lead to changes that would remind us of the changes undergone in individuals we once knew well but who subsequently spent years as alcohol abusers with resulting brain damage and personality changes, or indeed, the changes we see in Alzheimer's patients. If we could draw a line at some point in the progression of Alzheimer's and say "to here and no further", and we related that to a proportion of damage in the brain's ultrastructure, then I think we could come up with a measure for "fatal information loss".

2.) What would be good cryo-preservation if we wanted to devise a recovery method at this time?

I am not so good at pretending to be a superintelligence from the future and I tend to think how we could develop a project to achieve specific objectives right now. Where cryonics is concerned, my overriding interest in substrate-independent minds and whole brain emulation makes me less interested in any kind of "revival" from cryo through biological healing of tissues. I would aim squarely at recovery through whole brain emulation. So, for me the things to determine about cryopreservation are exactly those that we worry about for data extraction to achieve WBE.

I've argued before that structure is probably not enough. Functional

reference points are needed to turn the problem into a manageable collection of small system identification problems. Each reference point allows us to test our assumed parameters, to see if the subsystem produces the responses we would like it to produce. If that is absolutely necessary, then I should probably argue that proper cryopreservation requires either that:

a.) The brain can be revived (at least in slices) so that response measurements can be made. But preferably that,

b.) a very large number of high resolution response measurements should be taken even before the brain is cryopreserved. In other words, go through a mature data acquisition procedure for whole brain emulation, instead of simply cryopreserving.

It might be that if we had much better neuroscience understanding then we would know exactly what elements of structure and (protein)chemistry to preserve so that we could infer all of these details from morphological and chemical examinations of preserved brains. I do not think that we have this understanding at this time.

To me, this means that it is quite likely that some unrecoverable information that is quite important to reconstruction is presently lost during cryopreservation. If so, then the results of an attempted recovery would be – well, interesting. It should certainly be tested, because experiment is the best way to augment our lacking knowledge.

3.) Recovery by super-duper intelligence, regardless of lost information:

This reminds me of a recent thread that invoked "Quantum Archeology" to bring back anyone who had every lived due to the theoretical reversibliity of processes in a quantum mechanical universe. People seem to think that this means you could use an intelligent computation to sift through states of the universe (or a system, e.g. a cryopreserved brain) and determine the true prior states so that no information is every "lost".

Theoretically, perhaps… but it would probably be a good idea to listen to the cool talk by Dr. Ron Garret (http://youtu.be/dEaecUuEqfc) that I recently posted about. Among many other things, he addresses the matter of reversibility in quantum mechanics. Unfortunately, to reverse states, you actually need to carry out physical reversal, bringing particles back to their original locations, etc. In other words, it may be theoretically possible, and information may never be truly lost in the most absolute sense, but it is completely impractical – no matter how intelligent you are.

So, at the least we can state that this notion of recovery by superintelligence is not an absolute possibility regardless of the state of information loss. No doubt, a better understanding of neuroscience, a large amount of experience with the recovery of cryopreserved patients and computational aids would enable recoveries from conditions that would otherwise not be recovered from… but do those conditions correspond to the conditions that current cryopreservation can achieve? I have no idea.

4.) Personally successful recovery:

Our experiences, our sense of self, personal identity, all of those are the result of mental processing. The feelings they produce are internal to us. They are just as amenable to modification as anything else in a mechanistic process.
In short, it should be possible to engineer a sense of successful recovery in any recovered individual… no matter whether that individual actually resembles the one who entered cryopreservation. We can believe pretty much anything. So, from a personal perspective there is no need for any recovery ever to seem unsuccessful.
But… was that sensation really the aim? 😀
These fundamental questions about being, experience, self, etc. are of great interest as we begin to be able to deal more directly with them. Weird as such a solution may sound, I think it at least needed to be mentioned, for the sake of thinking outside the box.

So, my personal view of cryonics? Despite all of the uncertainty, it seems like a matter of personal priorities: Do I prefer to spend $50-$100 per month on an insurance policy for cryopreservation, or do I prefer to spend it on coffee and a restaurant dinner? I think the value proposition is worth the investment, for me.
I think it is too soon to draw conclusions where we would impose a responsibility on care givers to insure cryopreservation, because there are so many other things that care givers also have to make decisions about that are investments of effort and resources. All of those should appear on one's list of considerations and be prioritized in a manner that suits personal objectives. That care givers make decisions for children who cannot yet take charge of their own priority list is an issue that comes up often in many circumstances. It's worth reevaluating regularly.

4

Tags: cryonics

1 comment

0

Leave a message...

Discussion | Community | Share

Dr. Ridell • 5 days ago

I have decided to sign up with the cryonics institute because for 28,000 it is a relatively cheap gamble. Life insurance covers it and if it doesn't work I am barely worse off then someone who never signed up. To me its a no brain-er.

0 • Reply • Share ›

ALSO ON ALEXANDER KRUEL

What's this?

Why you should be afraid of friendly AI

2 comments • 16 days ago

dmytryl — I don't think you should be any more worried about SI than you would be worried about cold …

Richard's paradox

1 comment • 5 days ago

Duck — The errors is that you can enumerate an infinite list in this way. It is similar to the Godel diagon…

The Singularity Institute: How They Brainwash You

2 comments • 11 days ago

Alexander Kruel — The point about method 3 is that one could replace the subject of argument, in this …

Yvain on not taking ideas seriously

1 comment • 15 days ago

dmytryl — He's still endorsing people taking ideas seriously, if you read to the end. And calling it …

Comment feed Subscribe via email