# How to avoid steering blindly:
# The case for a robust repository of human knowledge

Jens C. Niemeyer*

*Institut für Astrophysik, Universität Göttingen, Germany*

(Dated: April 16, 2014)

## Abstract

Steering the future hinges on the availability of scientific and cultural data from the past. As humanity transitions into the digital age, global access to a condensed form of human knowledge becomes a realistic technological possibility and potentially a human right. At the same time, the risk of losing the vast majority of this information after a global disaster has never been greater. I argue that a collaborative effort to create a secure repository of human knowledge would not only protect humanity's cultural heritage for future generations, it could also define a minimum standard for the information that every human being should have a right to access. The basic requirements and challenges for creating the repository are discussed.

* jens.niemeyer@phys.uni-goettingen.de

## I. INTRODUCTION

Throughout human history, two elements have proved essential for the development of human culture: the capability to conserve and communicate previously gathered knowledge about humans, their interactions, and the world they live in, and the capacity to derive "useful" (in the sense of being benefitial for the long-term survival and development of the respective group of humans) decisions based on the available information. The basic assumption of this essay is that this observation will remain true for the way humanity will steer the future.

The second element, decision making, has been the responsibility of tribal leaders, monarchs, or elected officials using guidance provided by different types of expert committees (elders, spiritual advisors, scientists etc.) which, in turn, have used the best available sources and technologies at their disposal. Looking forward, great efforts are being made to employ simulations of complex social and economical interactions to advance the technical support for decision-making (e.g., the "Living Earth Simulator" [1]). While this is undoubtedly a central ingredient in humanity's capacity to steer the future, it will require major technological breakthroughs and its ultimate success may be fundamentally constrained by limitations in the predictiveness of models of complex systems far from equilibrium.

By comparison, the first element, i.e. conserving and making accessible the essential parts of human knowledge (where defining the metric for "essential" will be one of the key challenges) into the far future, has received much less attention. This is critical, because as humanity transitions into an age where information exists nearly exclusively in digital form, this knowledge is far more vulnerable to major disruptions of technological infrastructure than ever before. The risk of losing most our of cultural and scientific heritage after a global disaster (for instance, any event that interrupts electrical power supply on a global scale for more than a couple of weeks) is substantial unless precautions against this *digital amnesia* are taken.

On the other hand, digital technology also offers the opportunity to create a universally accessible, multi-dimensional and multi-resolution repository of human knowledge (henceforth simply referred to as the *repository*) with an enormous potential for research and education in "normal" times. As will be argued below, the internet already provides several promising services for information conservation and retrieval, but none of these satisfy all

of the central requirements of the envisioned repository. No fundamental breakthroughs are needed to begin this project, and the risk of digital amnesia adds a certain sense of urgency. The goal of this essay therefore is to demonstrate the potential benefits of a global knowledge repository (sec. II), to summarize the basic requirements derived from these intended purposes (sec. III), and to outline some ideas for its creation (sec. IV).

## II. BENEFITS OF THE REPOSITORY ON DIFFERENT TIME SCALES

The diverse benefits of a global knowledge repository can be demonstrated most clearly by viewing the future of humanity on different time scales. This will allow us to identify its key requirements specifications.

### Near future

Already at the present, research and education profit greatly from open online resources that cover different parts of the spectrum of human knowledge. Wikipedia is probably the most prominent example for a very broad but relatively shallow organization of knowledge (we might call it "horizontal") , allowing the efficient retrieval of information about an extremely wide range of topics, but being poorly suited for an in-depth study of any given subject starting from zero and going to the frontiers of current research ("vertical" organization). On the other hand, a wide variety of courses and tutorials ranging from short How-To clips on YouTube to postgraduate level Massively Open Online Courses (*MOOCs*) are available for the latter.

Ideally, a knowledge repository should allow flexible navigation in depth and breadth at an arbitrary level of detail ("multi-resolution") within the space of (appropriately linked) fields of knowledge ("multi-dimensional"). In other words, it should provide overlapping maps of the space of human knowledge with adjustable resolution, i.e. an atlas of knowledge space. Clearly, the atlas must continually evolve by ingesting new information and re-organizing existing correlations. This can only be achieved if the repository acquires a certain (and growing) degree of *autonomy* using artificial intelligence, combined with human supervision, for creating and updating maps.

The closest existing structure to the proposed repository is perhaps the *Internet Archive*

[2] operated by a non-profit organization that aims to provide permanent access to historical collections in digital format. However, the repository's purpose would go far beyond static preservation of data. In order to allow humanity to "bootstrap" its knowledge in the aftermath of a global disaster, great emphasis must be placed on i) the self-containedness and self-consistency of the information in the spirit of an enormous collection of cross-linked textbooks (starting with tutorials to read and write), and ii) the robustness of the data in the event of long-term disruptions of infrastructure (see below).

**Intermediate future**

On the scale of several decades, free access to essential knowledge must (and probably will) become a major political goal and may be considered a basic human right. The definition of "essential knowledge" obviously depends on the cultural and religious perspective. Nevertheless, it is crucial for the future of humanity that a consensus can be found on what information is considered essential for conservation into the far future, where "information" refers to facts and all the necessary correlations to understand them from scratch (the *maps* introduced above). At the same time, this body of knowledge can define the minimum level of information that every human being should be entitled to having access to. We should hence add *accessibility* to our list of key properties of the repository.

If the repository project becomes sufficiently visible, it can perhaps trigger a global political and cultural debate about the importance of access to information as a human right. This may indeed contribute substantially to help humanity steer the future.

**Far future**

Looking into the more distant future, humanity will face global disasters (cometary collisions, epidemics, nuclear wars etc.) with finite probability. Some of these will disrupt human society and its technological infrastructure for long periods of time. One of the main objectives of the knowledge repository must be the conservation of humanity's cultural and scientific heritage and the "re-booting" of human society after such catastrophic events.

This purpose of the repository is similar in spirit to the mission of the "Albertian Order of Leibowitz" in Walter M. Miller's classic novel [3][4]. However, my vision of the future and

of humanity's handling of scientific knowledge is considerably more positive than Miller's.

Currently, no widespread efforts are being made to protect digital resources against global disasters and to establish the means and procedures for extracting safeguarded digital information without an existing technological infrastructure. Facilities like, for instance, the *Barbarastollen* underground archive for the preservation of Germany's cultural heritage (or other national and international high-security archives) operate on the basis of microfilm stored at constant temperature and low humidity. New, digital information will most likely never exist in printed form and thus cannot be archived with these techniques even in principle.

The repository must therefore not only be *robust* against man-made or natural disasters, it must also provide the means for accessing and copying digital data without computers, data connections, or even electricity.

Speculating about the ultra-distant future, mankind may spread out inside the solar system or beyond. In this case, the repository would fulfill yet another purpose: the condensed human knowledge to be carried along wherever humans travel.

## III. FUNDAMENTAL REQUIREMENTS

We can now address each requirement in turn to discuss the main challenges for the development of the repository.

### Semi-autonomy

Long-term data archiving is already an active field of research in the information sciences. However, the envisioned knowledge repository is far more ambitious than any existing approaches since it intends to organize *all* forms of knowledge into an atlas of overlapping maps. Furthermore, none of this information can actually be archived into slow-access storage, as the continuous map updating process will require fast access to essentially all of the data.

The human brain is the most powerful known realization of a flexible storage and retrieval mechanism for multi-dimensional and multi-resolution facts and correlations. Creating and updating maps from a continuous data flow is similar to committing memories to long-term

memory in the human brain. It is therefore plausible to assume that brain research, such as the Human Brain Projects funded prominently in Europe and the US [5], may provide useful insights for the software design of the repository. More generally, further research in neurophysiology and artificial intelligence are needed to make the map-creating process partly autonomous in the long run to keep up with the ever-increasing amount of data.

While semi-autonomy of the repository for data ingestion and re-organization will be crucial on longer time scales, it is important to stress that it is not a pre-requisite for starting the project. Much of the initial structure of the repository can and must be constructed by human developers.

**Accessibility**

During times of a fully operational global technological infrastructure, access to digital information from nearly every place on the planet presents only practical, but no fundamental technological obstacles. The biggest challenge here might be the avoidance of a pronounced social gradient, i.e. access to digital information must be provided for all humans regardless of their social, geographical, or political status.

The repository project can immediately aim for a free, multi-device internet service. It would be desirable to keep it independent from corporate and political interests as far as possible.

In contrast, providing accessibility in the event of a global disaster presents a serious technological (and sociological) challenge for the project. The repository must feature a "bootstrapping mode" which allows the regrowth of knowledge from nearly zero. A few ideas will be sketched below, but clearly this question warrants a dedicated research initiative.

**Robustness**

Despite some efforts to build ultra-secure data storage facilities, it is probably safe to assume that most existing data centers would not survive a disaster on continental or global scales for more than several months. Stability of power supply and finite redundancy against hardware failure are among the most serious long-term problems.

On the other hand, not all of the data would have to be directly accessible for all times.

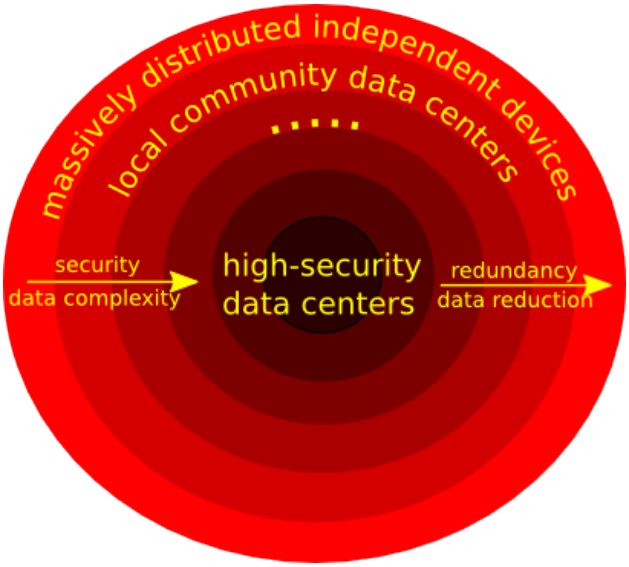


FIG. 1. Illustration of the onion-skin strategy for protecting the repository against global disasters. The outer layers should provide robust and redundant access to basic information (tutorials and instructions) which is then used to successively recover the inner layers.

Instead, the repository could allow for a hierarchy of resolution and dimensionality of the stored knowledge maps in a sort of onion-skin structure, as illustrated in fig. 1. The outermost layer of data, having the highest level of redundancy and post-crisis accessibility, would only need to contain the information needed to restore technology to the level where the next layer can be resurrected, and so on until all the necessary infrastructure has been recreated to access the entire data (which might be centuries later).

For example, a strategy that combines ultra-secure, centralized data centers with massively distributed independent devices (along the lines of the *One Laptop per Child* project [6], albeit with different goals and distribution strategies) might guarantee the intended degree of robustness and redundancy. In this case, the independent devices (e.g., solar powered, low-cost tablets) would represent the outermost layer of the onion skin, containing, for instance, simple tutorials for survival, basic social and technical infrastructure, as well as instructions on how to find and operate the next level storage facilities, which might be local data centers, etc. At the core of the onion, underground facilities designed to survive for centuries without maintenance (such as, e.g., [7]) would store the highest resolution knowledge maps. Thinking very far ahead, even a lunar outpost of the repository might be considered to be on the safe side.

## IV. OUTLOOK: HOW TO GET THERE

First steps toward the implementation of the repository can be taken immediately. In the beginning, a crowdsourcing approach might be most promising, supported by a coordinated initiative to attract research institutions to work on the fundamental design challenges regarding robustness and semi-autonomy described above. The goal of this initial phase would be the establishment of a global "repository collaboration" and first demonstrations of the feasibility of the project.

In the long run, however, the repository should strive for some sort of global political recognition that supports (if only symbolically) its relevance for the future of humanity and protects its existence and independence. This might be achieved, for instance, under the umbrella of UNESCO.

Ultimately, the protection and support of the repository may become one of humanity's most unifying goals. After all, our collective memory of all things discovered or created by mankind, of our stories, songs and ideas, have a great part in defining what it means to be human. We must begin to protect this heritage and guarantee that future generations have access to the information they need to steer the future with open eyes.

---


[1] M. Paolucci, D. Kossman, R. Conte, P. Lukowicz, P. Argyrakis, A. Blandford, G. Bonelli, S. Anderson, S. Freitas, B. Edmonds, *et al.*, The European Physical Journal Special Topics **214**(1), 77 (Nov 2012), `http://arxiv.org/abs/1304.1903`.

[2] `https://archive.org`.

[3] W. M. Miller Jr., *A Canticle for Leibowitz* (Bantam Spectra, 1984), ISBN 9780553273816.

[4] An alternative title for this essay might be "What would Leibowitz do in the age of ebooks?".

[5] E. R. Kandel, H. Markram, P. M. Matthews, R. Yuste, and C. Koch, Nat Rev Neurosci **14**(9), 659 (Aug 2013).

[6] `http://one.laptop.org`.

[7] `http://www.swissfortknox.com`.