

# Humans Must Help Humanity Steer Itself

Jeff Alstott

## Introduction

The future is important; we are going to spend a lot of time there. How, then, can we ensure that humanity will get the future “right”, that civilization will have a good trajectory instead of a bad one? There is no simple answer, because reality is not simple. In this essay, I will describe how human society is an enormous, self-steering system full of nonlinear internal dynamics, which constantly interacts with the Earth, the solar system, and the rest of the universe. Regardless of one’s particular goals for society, steering humanity towards a specific future is a task too complex for any one human to understand. No human can set the course for all of civilization. But what humans can do is coordinate so the larger system of humanity can steer itself. This coordination allows humans to process information and perform actions collectively, in ways that no individual could. Humanity’s performance relies on the efficacy of this coordination, and therefore improving the coordination system can help improve humanity’s future.

I will argue that the performance of humanity as a whole is a separate issue from the performance of its subcomponents. We could have great governments, great economies, great philosophies and great physics, but if a threat like global warming destroys us, civilization as a whole will have failed. I will argue that we must explicitly evaluate and improve global society as a whole, and will identify a powerful mechanism for identifying well-functioning systems: competition. However, this mechanism will not work at the level of humanity, because we only have one civilization. I will describe ways to compensate for this limitation by using competition between other, surrogate systems to identify coordination processes that will likely improve civilization. By identifying and implementing superior systems for coordinating action and information processing, humans improve humanity, and thus humanity’s ability to find further system improvements. In this way humans

can recursively improve and create an ever-evolving system of humanity.

## **The System of Humanity**

Self-steering systems exist in nature at many scales. A cell is the canonical simple example of a semi-closed system that interacts with its environment to pursue a non-trivial goal or trajectory. Cells, however, are composed of organelles and finer chemical processes with similar self-steering qualities. More obviously, cells can work together to create multicellular organisms, such as humans. The dynamics and goals of these larger organisms emerge from the actions of cells, but are different and separate from those of the cells. Humanity is an organism constituted of humans. Its behavior emerges from the actions of humans, but its dynamics are not human-scale, its problems are not human-shaped, and its course will not be steered by individual human minds.

## **Coordination**

An individual human is likely able to cook a hamburger, but is likely unable to single-handedly operate a hamburger restaurant chain. There are two problems for the individual: one of action and one of information processing. The problem of action is obvious; the individual cannot be at every restaurant simultaneously. But this problem can be overcome by using more than one human body, via delegation: simply tell multiple people what to do. This may work at a small scale, with a few people all performing the same task. Scale up, however, and the authoritarian restaurateur cannot follow the details of all the activity across all the restaurants, and the hamburger chain cannot function. This is the problem of information processing. Global civilization overcomes the problem of information processing by using more than one human brain. Humanity coordinates the activity of many individuals humans to work together to collect, process, and analyze information. Individuals integrate their efforts to create a gestalt analysis and a course of action. Coordination is frequent within well-run businesses and governments, but coordination also exists across organizational boundaries and between individuals. When one reads a news article created by dozens of writers, editors, and eye-witnesses, that is coordination of mental effort. When one reads an essay, such as this one, which is informed by articles, textbooks, teachers, conversations, and first-hand experience, that is also coordination of mental effort. Coordination of action and thought can be

very powerful, and accordingly it is a pervasive feature of human civilization.

## Competition

Countless methods and structures exist for coordinating human effort, but not every coordinated effort produces good analyses or actions. The preeminent mechanism for identifying effective coordination systems is competition. In competition, multiple systems individually try to perform some target action. Systems that perform the action well compared to others increase in some manner, while those that perform poorly decrease. These systems could be cells, companies, countries, and so on. The actions could be to make offspring, to make profits, or to make happy humans. High performing systems could increase in number or size, or simply persist longer than others. Performance can depend on the environment; a system may perform well in some environments, but not others. If the environment is noisy, a system will need to be robust to such variance in order to perform consistently. If the environment changes over time, consistent performance will require the system to take in information from the environment and adapt accordingly. Competition can be a powerful force for shaping systems, including creating systems that are robust to noise and adaptive to environmental changes. These are exactly the qualities society needs to have to thrive into the future.

Nature has used competitive forces to shape systems like organisms, and we make use of competition within society to shape such systems as companies, economies, and governments. Competition, however, cannot operate on civilization as a whole, because we have only one humanity. Improvements from competition require multiple competitors; the good performers are increased, and the bad are decreased or removed. If there is only one competitor, there is no mechanism to improve. If the single competitor performs sufficiently, it survives. If it does not, it dies and that is the end. We have only one global civilization, and so we cannot rely on competition to tell us if our society is one that performs well. The only competitive signal we will receive that our civilization was insufficient is when it ends. Then it will be too late.

## The Task for Humans

Our problem is to create a society that will succeed, without ever observing it fail. Humans must coordinate so that humanity can steer itself, and that self-steering must be robust to changes on Earth and in the larger universe. We must create this robust, coordinated, self-steering system of global society, and we must do it in one shot. Work has already begun on this task; we already have a civilization. Depending on one's goals for humanity, civilization today may be considered exemplary. But both society and the external environment are changing, so future performance may change. The perfect solution would be a societal system that is mathematically proven to persist and perform in all environments. Unfortunately, such a solution is likely impossible, perhaps even provably impossible, unless the performance goal is extinction. We humans may never guarantee a perfect future for humanity, but we can give it ever increasing odds of a good future through continuous effort. Our task is to constantly increase and apply our knowledge of societal engineering.

How then can we improve civilization? We cannot use competition to optimize it directly, as we have just one. We can compensate, however, by studying surrogates of society. These surrogate systems capture what we hope are the relevant details of humanity's needs and capabilities. These systems will compete to show the best performance and robustness, and the best of them can then be implemented as our own society. Possible surrogates are subcomponents of society, past societies, and simulated societies.

## Sub-Civilizations

The most obvious reservoirs of surrogate systems are from such fields as political science and economics, which specifically study the behavior of human systems. These fields have the advantage of working with empirical systems. Data collection has historically been difficult, but is getting increasingly easier. Most of this research, however, has examined individual countries or markets within global civilization, and does not explicitly study the behavior of civilization as a whole. For some features, this distinction does not matter; a global society functions much like a smaller one. Frequently, however, "more is different".<sup>1</sup> Large aggregates of states do indeed have different dynamics from those of smaller polities (e.g. the United Kingdom is not simply a large London). This is evidence that the larger global system is functionally distinct from its smaller subsystems,

just like the human body is functionally distinct from its smaller cells. We must recognize dynamics that are the same on the global level as on the smaller level and use the knowledge gained from studying individual countries and markets. But we must also recognize which features are different, and explicitly study the behavior of human civilization as a whole.

## **Past Civilizations**

A way to examine humanity as a complete system is to study past instances of civilization. Since global society has changed significantly over time, the different civilizational systems across history are quite varied, as are the environments in which they resided (e.g. Ice Age humanity vs. Internet Age humanity). Studying the past has the advantage of measuring empirical data from real global-scale civilizations. It can also reveal trend lines, such as civilization becoming more robust or more fragile. The difficulty, of course, is the availability of data. More data from the past is always being uncovered, but there are clear limits on how detailed reconstructions of the past can be.

## **Simulated Civilizations**

Both current and historical surrogate systems provide empirical data, but are limited to the data available. Simulations can create any civilization we want, at the expense of empirical grounding. Simulating an entire society is technologically demanding, but it is now possible to perform such simulations at a very coarse-grained level, and our capabilities are growing. Simulating many, diverse civilizations would increase the fodder for competition, increasing the likelihood of finding superior ways to coordinate humans. High performing coordination systems could then be simulated again in radically different environments, thereby testing their robustness. Simulations will not be perfect representations of reality; some abstractions must be made. The utility of simulation surrogates will depend on whether these abstractions have removed important features of reality, so these abstractions must be continuously rechecked and reconsidered.

## **Borrowing Processes from Other Systems**

Competing surrogate societies are how we will test if a coordination system performs well. But we can look for additional inspiration from other self-steering systems that are apparently unlike societies. Research fields like biology, statistical physics and engineering have found mechanisms

that are clearly relevant for society, such as mutation, feedback loops, and modularity. Even the very process of “lifting ideas from other systems” is found in other systems: bacteria transfer genes between themselves and other organisms, and humans merge and recombine technologies to make new ones. Identifying coordination processes in other self-steering systems would allow for seeding simulated societies with these processes, to observe how they perform.

## Humans Creating Humanity

The goal is to identify and implement better societal systems, so that humanity can better take in, process, and act on information. But we are part of the social system. As we coordinate humans better, we will be better coordinated. As we develop and implement better methods for humanity to effectively digest and act on information, our own efforts will become more productive, because we will process the information better. There is thus a deep element of recursion to this process, whereby humanity grows into ever stronger structures. In order for this recursion to take place, however, better societal systems must not just be discovered, but also implemented.

This growth process has already begun: Essays like this one are submissions of candidate ideas for how to better construct society, and contests like this one are a method for these ideas to compete. What we learn here, if implemented, will help us to make new ideas and new selection methods, leading to yet further improvements in the coordination of human action and information processing. We are at this very moment participating in humanity constructing itself, steering its way towards an ever brighter future.

## References

<sup>1</sup> P. W. Anderson. More Is Different. *Science*, 177(4047):393–396, 1972.