Cliodynamics: The Journal of Quantitative History and Cultural Evolution

The Computational Power of a Human Society: a New Model of Social Evolution

David H. Wolpert[1], Kyle Harper[2]
[1]*Santa Fe Institute*
[2]*University of Oklahoma*

Social evolutionary theory seeks to explain increases in the scale and complexity of human societies, from origins to present. Over the course of the twentieth century, social evolutionary theory largely fell out of favor as a way of investigating human history, just when advances in complex systems science and computer science saw the emergence of powerful new conceptions of complex systems, and in particular new methods of measuring complexity. We propose that these advances in our understanding of complex systems and computer science should be brought to bear on our investigations into human history. To that end, we present a new framework for modeling how human societies co-evolve with their biotic environments, recognizing that both a society and its environment are computers. This leads us to model the dynamics of each of those two systems using the same, new kind of computational machine, which we define here. For simplicity, we construe a society as a set of interacting occupations and technologies. Similarly, under such a model, a biotic environment is a set of interacting distinct ecological and environmental processes. This provides novel ways to characterize social complexity, which we hope will cast new light on the archaeological and historical records. Our framework also provides a natural way to formalize both the energetic (thermodynamic) costs required by a society as it runs, and the ways it can extract thermodynamic resources from the environment in order to pay for those costs — and perhaps to grow with any left-over resources.

Introduction

How and why have human societies grown in scale and complexity over time? When H. *sapiens* emerged around 300, 000 years ago, humans lived in small bands of hunter-gatherers; energy harvest per capita was of the order of bodily metabolism (i.e. a small multiple of 2000 kC/per capita/per day); global population

Corresponding author's e-mail: david.h.wolpert@gmail.com

Citation: Wolpert, David H. and Kyle Harper. 2025. The Computational Power of a Human Society: a New Model of Social Evolution. *Cliodynamics* 16 (1): 69–122.

was low ($\sim 10^4$); and social roles were relatively unspecialized. The first humans had technology (fire, stone tools) and impressive cognitive abilities. Today, however, with essentially identical cognitive potential, H. *sapiens* numbers 8 billion individuals; we consume $\sim 10^{4-5}$ kC/per capita per day; and we inhabit globe-spanning social networks that facilitate extreme specialization (Morris, 2010; 2013).

Social evolutionary theories attempt to explain this growth over time in universal cross-cultural terms (Carneiro, 1967; Sanderson, 1999). Here, we propose a new framework for social evolutionary history grounded in the theory of computation. Our framework offers distinct advantages. First, by applying the formalisms of computer science theory in the context of human social systems, we can benefit from the rigor of this advanced body of thought about the nature of complexity; a sounder definition of human social complexity might be informed by theories of algorithmic and / or computational complexity. Two, it opens the possibility of a model of social complexity compatible with physics and biology, a longstanding goal of "big history." Third, given the rich and growing scholarship on the thermodynamics of computation, we can argue for the importance of coupled increases in energy harvesting and information processing as a key factor in human success. Both free energy harvesting and differentiation/specialization have a venerable place in the history of efforts to measure and model social complexity, so our framework is not incompatible with much existing empirical and conceptual work. The result might be a richer model of the dynamics of growth, seen as the interplay of information processing and energy harvesting.

Our idea is motivated by the success of computation-centered approaches in biology. Already Erwin Schroedinger, several years before the discovery of the DNA molecule, recognized that living systems must maintain themselves away from thermal equilibrium by consuming free energy and processing information (following an "elaborate code-script," as he put it [Schrodinger, 1944]). He also intuited the necessity of an "aperiodic crystal" for the storage of biological information that guides this process. In addition, the theory of "Major Evolutionary Transitions" (METs [Smith, 1999; Jablonka and Lamb, 2006]) in the history of life provides an important paradigm in which radical changes in biological information storage and processing led to major leaps in the complexity of the biosphere. More generally, there have recently been approaches to theoretical biology grounded in computer science, which are providing new insights into what life is and how it operates [Jost, 2020; Al-Hashimi, 2023]. The modern view is that organisms and biological systems can be viewed as essentially computational machines, which process inputs to produce outputs. The net result of that processing is forms of biological organization that reliably transmit information to future generations.

In light of this work in evolutionary biology, we view the emergence of the social systems of modern humans in the Holocene as an ongoing MET. In this current MET, human culture represents a new medium, complementary to DNA, for the accumulation and transmission of information about the external environment. A human society uses information, acquired and stored over the current and previous generations, to process inputs from its environment and to produce outputs, which are actions that it takes on that environment. In this sense, a human society is fundamentally a computer. We describe social evolution as the growth of humanity's collective metabolic-computational capacities. We are pursuing Schroedinger's observation that "living matter, while not eluding the 'laws of physics' as established up to date, is likely to involve 'other laws of physics' hitherto unknown, which, however, once they have been revealed, will form just as integral a part of this science as the former." Those laws, we believe, require accounting for increasing complexity via the interplay of computation and metabolism — on scales ranging from cells to societies.

What is the "information" stored in a human society and how is it used to perform "computations" on the environment? We think it is intuitive that this information is represented by the algorithmic instructions that the society runs in order to propagate itself into the future. More concretely, we propose that computation in social systems is the interaction of social agents with each other and with the environment to harvest energy, to reshape the environment, to reproduce, and to acquire new information. Over time, human societies have grown in scale and complexity because (to borrow terms from the economics literature) the collective "stock of knowledge" or "stock of ideas" stored in human culture has grown.

In essence, we propose to formalize the stock of knowledge as "information." In this view, ultimately it is not our sheer cognitive ability, nor our sociality (both of which are important to human success), that accounts for the apparent exceptionalism of H. *sapiens* in the stream of evolution. Rather it is our use of symbolic language, computationally powerful enough to support recursion and selfreference. Sophisticated verbal language, later amplified by information technology like writing and digital computers, allows the accumulation of information over time in culture (Hilbert, Lopez 2011). Culture, here, would represent the totality of the rules, recipes, blueprints, and instructions operated by a society.

It has been famously difficult to operationalize measurements of complexity, differentiation, and information processing in human societies. While our framework will not suggest any easy answers to this challenge, we will offer a practical strategy for combining theory, model, and measurement. We suggest occupational specialization as a proxy for specifying the information stored and processed in human societies. Since Adam Smith, specialization has been acknowledged as a main feature of economic growth. In Smith's words, "It is the great multiplication of the productions of all the different arts, in consequence of the division of labor, which occasions, in a well governed society, that universal opulence which extends itself to the lowest ranks of the people" (Smith, 1776). An occupation is like an ensemble of algorithms, a set of programs for how to interact with other agents and ultimately the external environment. An occupation is in essence a work tape, encoding instructions for performing computations. Occupational specialization has dramatically increased over the course of human history, as societies have grown and learned to perform a greater variety of operations on the environment. We propose that occupational specialization offers a practical, coarse-grained way to define and measure social complexity in computational terms.

We think that ultimately models of information and computation might also gain by greater dialogue with the social sciences, which present important realworld applications (but resist the sort of clean mathematization favored in textbooks on information theory and computation). If we are right, the road ahead is arduous. But because information processing is so fundamental to the universe, and to the nature of complexity, it is worth trying to reimagine the evolution of human societies as a story in which computation is central.

Where we are and where we want to be

In conventional historical social sciences, or some climate change time-series analysis, or some time-series analysis of epidemics, or even in vast swathes of evolutionary biology and allometric scaling theory, one can write down an extremely flexible set of parameterized equations for the dynamics of the variables that one is interested in (and can measure). Then one fits the parameters of those equations to a dataset of observations, to infer the dynamic "laws" governing the system. At root, such an approach is *phenomenological*, considering observations of how various aspects of the systems (and of human societies in particular) evolve over time, and using statistical analyses to provide insight into this phenomenology.

Needless to say, it will continue to be extremely useful to pursue such approaches further, tailored to investigating the dynamics of human societies ranging across pre-history (early Holocene) through the modern era. (Arguably, that is precisely the goal of the ongoing research in the Seshat project [Turchin et al., 2018; Sipper, 1998]).

However, such a phenomenological approach cannot give us any insight into the other METs not involving human societies; most if not all deeper understanding arising from applying this approach would be one-off, not directly transferable to different situations. And perhaps most importantly, there is no sense in which either information theory or computation modeling per se can play a role in this approach. Under such a phenomenological approach, the only role that the computation performed by a human society might play is to provide a motivation for a handful of definitions of "complexity characteristics" (to use the term in the Seshat project, which is distinct from the use of the term in computer science theory). Quantifying systems in terms of such complexity characteristics is analogous to characterizing a modern digital computer into three levels: strong, moderate or weak. It does not say anything specific about what computation is actually being performed on that digital computer (much less about what it would mean for a computer to jump from one level to another).

While we think it will be productive to pursue the conventional, phenomenological approach, we also hope to make progress on a fundamentally more challenging modeling exercise.

Nobody at present has a clear idea of how best to model computational systems outside of the purely mathematical, non-physical structures considered in CS theory or their direct translations into statistical physics models. (See [Wolpert, 2019] for the former and see [Wolpert, Korbel, Lynn et al., 2023] for further discussion of the latter point.) But at the same time, with a few exceptions focusing on biology (work by Gershenson, Zenil and Chaitin in particular comes to mind), there has never been a concerted effort to try to figure out just how to do that. Certainly it has never been extensively pursued in the context of deep history.

This essay maps out how such a project might be envisioned. In the next section (Section 2), we briefly discuss existing definitions of social complexity and introduce our own approach, grounded in computational terms. Next (Section 3), motivated by considerations of what data sets are available or potentially available, we propose a novel proxy for the "computational complexity" of human societies over time, taking occupational diversity as a metric. In the final five sections (4-8), we introduce a first-principles model of human societies as computers and suggest how we might capture the dynamics of social evolution, considered as the growth of computational-metabolic capacities.

We emphasize that the goals of this exploratory paper are to stimulate conversation about a research agenda rather than provide conclusive answers. We would expect pursuing such research might revise or fill out (or replace) the ideas sketched in preliminary form here.

Framing Social Complexity

What is social complexity? A complete discussion lies beyond the scope of this essay, but we do wish to retrace some steps where social theory could have taken a different path in the 20th century. We will also introduce and discuss a new framework for considering human social complexity. The hope is that a formal definition of social complexity, along with the associated framework that involves both computation and the harvesting of free energy by human societies, could address the deep mysteries of the growth of human societies in a richer way. To begin, we observe that two crucial intellectual developments have unfolded since the middle of the twentieth century, which are rarely observed in conjunction. First, evolutionary models of human history went mostly out of favor. This happened in part because the historical sciences turned away from hopes for unity with the natural sciences (Harper, 2013; Morris, 2022). Second, theoretical computer science (CS) took shape as a scientific field, with major advances in the definition and understanding of complexity (Mitchell, 209; Mokyr, 1990; Li and P, 2018; Arora and Barak, 2009). While there has been important cross-fertilization between CS, physics, and biology, to date there has been little overlap between the study of complex systems from a computational perspective and the evolutionary study of human societies. We will argue that social complexity is a quantity proportionate to a society's computational power.

A brief intellectual history of social evolution

In large part, this lack of contact between these fields is due to the configuration of academic disciplines. When the modern research university took shape in the late 19th century, an enduring division of labor took hold in the way the past is allotted to different fields (Smail, 2005). The pre-human past was allotted to scientists (e.g. geologists, paleontologists, and biologists). The prehistoric human past was assigned to anthropologists, who would specialize in the study of "primitive" societies either through ethnography or archaeology. Meanwhile, the last few thousand years – the period of governments and writing – became the domain of historians. This arbitrary division still affects the study of social evolution.

Stadial models of social evolution were developed already in the Enlightenment - by John Millar, Adam Smith, and the Marquis de Condorcet, among others. Herbert Spencer linked stadial history to biological evolution (and highlighted the importance of differentiation as an index of complexity). In its early decades, anthropology was dominated by social-evolutionary models (Tylor, 1881; Morgan, 1877), which proposed a progression from simple or primitive society to more advanced civilization. While value-laden models of "progress" were adopted by early anthropologists such as E.B. Tylor and Lewis Henry Morgan, subsequent generations of anthropologists reacted against existing stadial models of social evolution. There were, to be sure, western imperialist biases baked into this normative, progressive model of social evolution, which was often influenced by crude social Darwinism (Kuper, 1988). In reaction, Franz Boas and his influential students proposed relativist definitions of culture. Stadial models, ideas of progress, and evolutionary models came to seem hopelessly tainted with cultural prejudice.

The anti-evolutionary stance was then reinforced by the hermeneutic methods that became predominant in anthropology and the humanities more generally. It would be hard to overstate the influence of the so-called "linguistic turn" in professional academic fields such as anthropology and history. The idea that the study of the humanities was a pursuit apart from the study of nature was not new (at the beginning of the century, Wilhelm Dilthey had argued for methods specific to Geisteswissenschaften apart from the Naturwissenschaften). But the linguistic turn cemented the idea that the humanities were concerned with the interpretation of each culture in all its particularity. In this view, developed in the work of Clifford Geertz, human culture is a web of meaning, and the objective of anthropology or history is to interpret culture – to understand rather than to explain sequences of cultural development over time (Geertz, 1973). Qualitative methods were ascendant over quantitative ones, while paradigms drawing inspiration from the natural sciences were branded as "reductionist" (Harper, 2013).

As has been pointed out ([Morris, 2010], among others), this work criticizing social evolution does nothing to invalidate the possibility of an objective and neutral measurement of social complexity. A more complex society is not intrinsically "higher" or "better." But it is more complex (as objectively, we would argue, as one algorithm or computational problem can be more complex than another, without any necessary moral hierarchy). While research on the dynamics of social complexity has survived in pockets of anthropology, in the discipline of academic history, formal, mathematical frameworks were banished with almost religious zeal. Over the second half of the twentieth century, academic history, as a field, came to prefer qualitative over quantitative methods, to reject scientific models of explanation (on the ground that they are reductionist), and to focus on particularities over generalities. As Ian Morris has noted, whereas most social-science fields maintain a healthy tension between interpretive and explanatory approaches, the commitment of historians "to humanistic and particularistic questions and methods verged on monomaniacal" (Morris, 2013).

As a lingering consequence of the 19th-century division of labor, anthropologists still tend to study relatively small-scale, simple societies, and this is especially true of anthropologists who are most open to evolutionary thinking. "Cultural evolutionists" in contemporary anthropology use rigorous evolutionary models to help illuminate the ways in which culture is adaptive and might be transmitted; in particular, "cumulative cultural evolution" helps to explain how foraging societies can solve complicated problems via the accumulation and transmission of knowledge, stored in culture (Bettencourt, 2020; Henrich, 2004; Smaldino and Richerson, 2013). Historians, by contrast, mostly study the last few thousand years of the past, and since evolutionary models are effectively taboo, there are few evolutionary models that try to account for the spectacular growth of human societies in the late Holocene. The net result is that evolutionary models of human history have been nurtured mostly on the margins of anthropology, or within other fields altogether (such as economic history, or the "big history" which is mostly written by physicists and biologists), or among dissidents like Morris.

Thus, someone naively interested in how human societies have grown in scale and complexity would be surprised to learn that there does not exist a terribly robust literature on what social complexity is or how we have achieved it. There are of course exceptions. Even as Boasian relativism was ascendant, the anthropologist Leslie White sought to reformulate an evolutionary model on more neutral grounds, centered on energy (his model is focused on the following "equation": culture = energy times technology [Leslie, 1959]). Others have adopted the approach of Elman Service, who developed a definition of social complexity based on "levels of hierarchy," (e.g. band, tribe, chiefdom, state [Service, 1971]). And over the years, there have always been some (Raoull Naroll, Robert Carneiro, George Murdock, etc.) who have kept alive efforts to define and quantify social complexity, sometimes blending differentiation/specialization, hierarchy, and scale (e.g. the size of the largest settlement [Naroll, 1956; Carneiro, 1967; ck, 1997; Denton, 2004; Donaldson-matasci et al., 2010; Gedeon, 2018; Murdock, Provost, 1973]).

The paradigm of Big History addresses the question of social complexity and very often focuses on energy capture and exchange (Christian, 2004, Spier, 2015). However, it has so far not deeply engaged with informational or computational measures of system behavior, or models of dynamics grounded in information theory and / or CS theory (though see [Christian, 2017] for thoughts in this direction). A parallel effort is represented by cliodynamics. Cliodynamics has inspired an ambitious research agenda around social complexity, notably the Seshat project (Turchin et al., 2018). While most of its practitioners are amenable to evolutionary models, cliodynamics has so far not engaged extensively with CS theory or information theory. Most work in this vein, moreover, has so far focused on the development of agrarian societies; the Seshat data go as late as 1900 in some cases, but it is fair to say that the dataset is not designed to offer insights into the dramatic increases in the complexity characteristics of modern industrial and post-industrial societies. In many world regions, maximal complexity was reached hundreds or thousands of years ago.

Economics has arguably been the most successful branch of social science in addressing big-picture questions about growth and development quantitatively, and here there are many potential points of engagement. First, there is the literature on economic complexity, which, while not historical in its orientation, provides rich models and datasets for contemporary economies (Hausmann, Hidalgo and Bustos, 2013). More historically, the "new institutional economics" highlights the importance of (formal and informal) rules in shaping exchange and production (Douglass, 1990). We would simply translate these insights into the terms of CS theory, to say that institutions or rules are part of the way that agents interact to process information. Much work in growth theory has underscored the importance of ideas and innovations in increasing productivity. Joel Mokyr, for instance, traces the breakthrough to modern growth back to the spread of practical, empirical science (Mokyr, 1990). "Unified growth theory" focuses on the dynamical feedbacks that encourage the formation of human capital and leads to the economic-demographic transition from a Malthusian regime to a regime of growth (Galor, 2011). Paul Romer won a Nobel Prize for his seminal contribution to endogenous growth theory, the "nonrivalry of ideas" (Romer, 1990). Unlike other goods that contribute to productivity, ideas are not diminished in being replicated. The Haber-Bosch process, for example, is based on a chemical reaction for synthesizing reactive nitrogen; it has contributed as much or more to human well-being as any innovation, and yet it is a simple recipe or idea that once discovered could be disseminated and replicated easily (relative to its original discovery). Most real growth can be attributed to increases in such a stock of knowledge or ideas. We see our proposal as compatible with these versions of growth theory, but we seek to formalize the nebulous "stock of knowledge" or "ideas" as information.

Computational approaches to complexity in biology

One consequence of the separation of the historical sciences from the natural sciences is that the study of human social evolution has been largely out of touch with developments in the field of complex systems science. Social evolutionary models have always tended to operate with an informal version of what "complexity" is - taking a "you know it when you see it" approach and sticking to proxies like specialization, hierarchy, and scale. This really bears underscoring. Even among those who think explicitly most about social complexity, there is very little effort to define complexity formally, much less in terms that are independent of the domain of human societies. But from the 1970s on, important developments have helped clarify what complexity means. A complex system is characterized by the interaction of parts, such that the dynamical properties of the system are emergent from the interaction itself rather than reducible to the properties of the system's subcomponents. It should be emphasized that there is no consensus definition or model of complexity. However, much progress has focused on a cluster of related ideas centered around algorithmic complexity (sometimes called "Kolmogorov complexity," after the Soviet scholar Andrey Kolmogorov) and computational complexity.

First, CS theory developed the concept of the "algorithmic (Kolmogorov) complexity of an object (represented as a bit string)", defining it as the shortest program one could write in some fixed programming language like Python that would reproduce that object and then halt. Algorithmic complexity provides a way to formally capture how much effort needs to be expended in providing instructions to a computer for how to construct an object.

Complementing this kind of complexity, computational complexity considers how difficult the actual construction process is (Moore andMertens, 2011; Motchell, 2009; Li and P, 2008; Arora and Barak, 2009). Whereas algorithmic complexity concerns an individual object (bit string), computational complexity involves problems, which are defined as (usually) an infinite set of questions of the form, "What is the optimal value of {...} in situation {...}?", using some pre-defined notion of "optimal". A famous example is the "traveling salesman" problem (TSP), defined as all instances of the question, "What is the shortest possible route a salesman could follow that connects a set of cities at the following coordinates?" (Arora and Barak, 2009; Mitchell, 2009; Moore and Mertens, 2011; Spiser, 1996). Computational complexity asks, "As the number of bits it takes to specify the situation grows, how fast must the resources required by a program to answer the associated question increase?" As an example, if we take "resources" to mean the time the program takes, computational complexity would consider questions like, "As I increase the number of cities in a TSP question, how much extra time must be taken by any program to answer the question correctly?"

While these new ideas of complexity have left little (if any) impact on the study of human social evolution, they have had a deep impact in evolutionary biology. The discovery of DNA in the 1950s confirmed Schroedinger's intuition, and the flowering of genetics has underscored that information is central to life. Biological inheritance is the transmission of information about the environment and how to do computations in the environment in a molecular substrate. In the words of Richard Dawkins, "what lies at the heart of every living thing is not a fire, not warm breath, not a spark of life. . . it is information, words, instructions" (Dawkins, 1996). Selection ultimately acts on information (Adami, 2012). The potential of information-theoretic notions of complexity inspired Eors Szathmary and John Maynard Smith's influential model of METs (Szathmary and Smith, 1995; Szathmary, 2015). These transitions are characterized by changes in the way that biological information is stored, exchanged, and transmitted. The canonical METs includes the emergence of life, protocells, genes, the eukaryotic cell, plastids, multicellularity, eusociality, and human language. The precise list is debated and debatable, but the concept is invaluable, foregrounding the interplay of metabolic and computational capacities in the history of life. The MET framework, with information processing at its center, undergirds most serious attempts to think about complexity in biological evolution.

The eukaryotic cell - arguably the most important development in evolution after the emergence of life itself, and the basis for all complex organisms - is a perfect example. In its mature form, the eukaryotic cell combines revolutionary increases in information storage and processing with much more powerful metabolic systems (epitomized by the mitochondrion). This is illustrated by the leap in power harnessed per unit of information in the transition from prokaryotes to eukaryotes (Fig. 1).

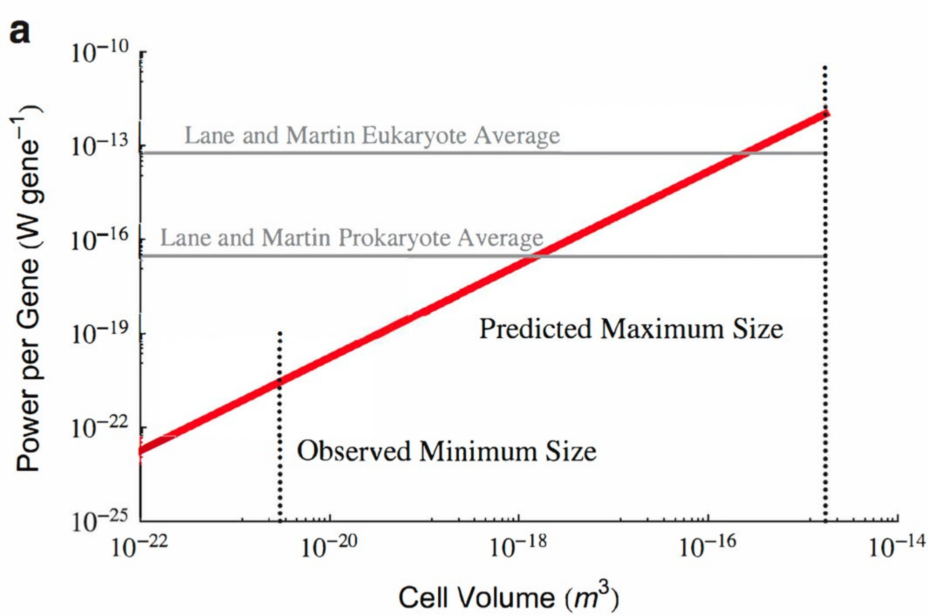

Figure 1. Scaling of metabolic rate to cell volume. The y-axis represents the power harnessed per unit of information, in watts per gene, with prokaryote and eukaryote averages in grey lines. Figure from (Kempes et al., 2016).

The computation-centered view of biology helps to explain what is, from a physics perspective, the central conundrum of life: how, in a universe dominated by the Second Law of Thermodynamics, living systems reduce local entropy, creating order and organization (Goldenfeld and Woese, 2011). The growth of the biosphere is driven by coupled increases in the metabolic-computational capacities of living systems; evolution is a kind of “Maxwell’s demon” sieving for information that can exploit free energy to reproduce itself (Krakauer, 2011). The information stored in living systems is information about the external environment. It is “how-to” information, ensembles of algorithms for how to self-assemble, how to observe the external environment, how to harvest free energy, how to act in and modify the environment, how to interact with other agents (conspecifics, predators, prey, etc.), and how to reproduce. For example, oxygenic photosynthesis, from this perspective, is an algorithm, a supremely successful set of step-by-step instructions for operating a chemical transformation that turns carbon dioxide, water, and sunlight into sugar. The “idea” of photosynthesis can be translated into chemical or even binary symbols ($6CO_2 + 6H_2O \rightarrow C_6H_{12}O_6 + 6O_{24}$). Of course in plants this program is operated by a network of organelles, enzymes, and genetic material, but it may indeed offer more “bang per bit” than any other genetic algorithm on Earth.

Needless to say, informational and computational approaches to biology are not always appropriate or useful at a given level of investigation. But they are truly fundamental, related to the deep connection between the rules of the living world and the physical universe. They help to address the biggest questions about the expansion of life. The “other laws of physics” that Schroedinger foresaw are the laws of complexity - of emergence and interaction, shaped and constrained by the interplay of metabolism and computation. We propose that these other laws of physics be extended to human social evolution.

Just as the long-run growth of the biosphere (measured in biomass, energy harvesting and exchange, and complexity) is a central question of macroevolution, so the long-run growth of human societies is the central question of human social evolution. We agree with Maynard Smith and Szathmary that the emergence of humans marks a MET and that language is central to this transition. Yet it is not just the emergence of our species that requires exploration in these terms. The last 300K years of human history can be considered an ongoing MET. We take growth - in scale, metabolism, and computational complexity - as the proper explanandum social evolutionary theory.

In the sequel we present a mathematical definition of a new kind of computational machine, and our framework for modeling the co-evolution of a society and its environment as the joint dynamics of two of these machines. We had several goals when constructing this framework, which in some senses were at crosspurposes. Foremost, we want the framework to be well-suited to investigating the issues described. In addition though we want it to be possible (or at least conceivable) to gain insight into its behavior using purely mathematical analysis, e.g., by appropriately extending CS theory. We also tried to ensure that the framework can be gainfully applied to investigate multiple METs, not just the current one. These goals drove us to try to formulate as simple a framework as possible, i.e., as "coarse-grained" a framework as we could. On the other hand, another goal we had was that it be possible to readily tailor the framework to the special case of the development of human societies. In particular, we want it to be possible to tailor the framework to incorporate the kinds of data sets we either now have or will in the near future. This drove us to try to formulate as rich a framework as possible, i.e., as "fine-grained" a framework as we could. Trying to meet these opposing goals led us to a framework that is "moderate-grained", in the same sense as the frameworks in (Arroyo et al., 2022; Bettencourt, 2013; 2020; West et al., 1997; 2001).

Measuring Computational Complexity in Human Societies: A New Proxy

There is a straightforward reason why energy harvesting has been an appealing index of social evolution. Energy harvesting can be measured in objective, quantifiable units (calories or joules) which are physically constant over time and space, permitting direct comparison: 500 kC of antelope meat on the African savanna 100Kya can be directly compared with 500kC of quinoa salad in a Santa Fe bistro in 2023. But what is measurable is not necessarily what is important - or in this case, is not the whole story. Even though information is also a fundamentally physical entity, quantifiable in bits, it is much more challenging in practice to measure information storage and processing outside of simple systems or mechanical computing devices. It is far from clear how much computation is occurring even in simple cellular operations that can be directly observed in laboratory conditions. Measuring information and computation in human societies is a daunting challenge. For past societies, yet more challenging still.

Various proxies for information storage and processing in human societies could be envisioned, and it is ultimately a question of both practical and theoretical considerations how best to approach the flow of information over time in human societies. For instance, we might just measure the scale of the "datasphere," the number of bits stored in extra-bodily media such as books and memory chips (in 2025, it is estimated globally to be 175 zetabytes). While this is not uninteresting, obviously many of these bits are inconsequential - think cat videos rather than the Haber-Bosch process. To put it formally, they are "syntactic" (independent of meaning) rather than "semantic" (meaningful, in the sense of causally effective: see [Kolshinsky and Wolpert, 2018]). In the pithy formulation of Artemy Kolchinsky and David Wolpert, they have little "bang per bit." A better measure has been proposed by Ian Morris, whose Social Development Index includes an Information Technology score (which takes into account literacy rates, writing, information storage, and communication technology) on a scale from 1-100 (Morris, 2010).

While this is still admittedly "crude" and not paired to any formal theory of information, it has the distinct virtue of making a start of things.

Our goal is to develop a proxy that best aligns with the theoretical framework we are proposing. In Section 4, we introduce a formal model of human societies as hierarchies of interacting computational machines, where each separate machine can be identified in many different ways as components of real human societies. As an example, we could identify each machine with a separate human. This would be particularly appropriate for analyzing small bands of hunter-gatherers. Another possibility, appropriate for analyzing larger Neolithic societies, would be to identify each machine as a separate city in a region, e.g., ancient Mesopotamia. Alternatively, focusing within a single such city, we could identify each machine as a separate occupation, and / or technology. Moving up to the present day, we could identify each machine as a separate firm in a modern state's economy.

Note that in all of these examples, the precise set of machines will vary over time, and one of the major "tasks" of the overall system of interacting machines is to extract energy from the physical environment, both to feed the members of the associated society, and for the society to use as goes about its activities — a key one of which is extracting yet more energy from the environment.

3.1 Occupations as (computing) machines. To choose among these (and other) choices for how to identify the separate machines in a (computational model of a) human society we also need to acquiesce to the reality of what data are available or potentially available. Accordingly, as a novel proxy for the computational power of a human society, we propose that occupational specialization is promising. (See (Hausmann, Hidalgo and Bustos, 2013), who in a similar spirit suggest thinking of the knowledge or expertise employed by an individual in terms of "personbytes"). An occupation is a job or a profession. It is a social role centered on labor, irrespective of the extent to which that labor is implicated in market exchange. Occupations can be more or less specialized and require higher or lower degrees of skill and/or training (i.e. from an information-theoretic perspective, not all occupations are equally complex, which means that it would be desirable to do more than simply count occupations to account for the computation going on inside an occupation). Since Adam Smith, the division of labor has been recognized as a hallmark of economic growth. Smith emphasized that the division of labor is limited by the extent of the market, but as Gary Becker and Kevin Murphy long ago observed, specialization is also constrained by coordination costs and — crucially — the total stock of knowledge (Becker and Murphy, 1992). The last is a particularly profound insight, although remarkably it has had little influence. We wish to build on it by leveraging the fact that occupational specialization therefore reflects the stock of knowledge. *Ceteris paribus*, the occupations that exist in a given society embody the totality of what a society knows how to do. We also immediately confess that one could make this proxy arbitrarily more sophisticated - and complicated - by considering technologies as computing machines too, since the integration of technology into human societies both replaces occupations and creates new ones. For the first step, we set this issue to one side, a concession to practicality rather than a claim that it is ultimately appropriate.

Occupational specialization reflects both diversification (the expansion of the number of things that a society knows how to do) and division of labor (the finergrained separation of functions and tasks in an economy, as in Smith's pin factory). Occupational specialization offers the potential to get at the total information stored in a culture and the processing of that information, i.e. computation. The importance that we assign to occupational specialization as an index is compatible with theories of economic growth that emphasize the total stock of knowledge or ideas (ultimately turned into innovations) as the most fundamental cause of long-term gains in productivity (e.g. [Easterlin, 1998; Jones, 2002; Koyama and Rubin, 2022]).

In this view, advances in the Scientific Revolution (e.g. Newtonian mechanics) were ultimately translated by engineers and entrepreneurs into useful productivityenhancing technological innovations like steam engines (and there are extensive literatures on how discoveries become innovations, who mediates this process, the extent to which innovation is evolutionary, etc.). The Industrial Revolution was transformational not because coal was suddenly available (it was always there), but rather because societies learned how to exploit fossil energy by creating blueprints of steam engines that transformed the mining, textile, manufacturing, and transportation sectors. The Industrial Revolution was, in our terms, a revolution that coupled increases in information processing and energy harvesting; for the first time in our history, these increases were of a speed and magnitude to permit escape from the Malthusian trap. The Second Industrial Revolution (ca. 1880-1930) was spurred by basic discoveries in chemistry and electromagnetism, which eventually produced massive innovation (electricity, light bulbs, telecommunications, internal combustion engines, fertilizer synthesis, polymers, etc.) that spread to scale worldwide over the 20th century. In recent decades, the mass production of microprocessors has been arguably the most important innovation.

We imagine that the aggregate stock of information encompasses the basic science (e.g. Maxwell's equations) and practical implementation (silicon is a good material for circuits); this aggregate stock of information is then reflected in the composition of an economy's occupations (e.g. research chemist, software programmer). As Hayek recognized, perhaps the most amazing thing that economies do is to coordinate vast amounts of knowledge, even though individual agents themselves only have access to a tiny proportion of that knowledge (Hausmann, Hidalgo and Bustos, 2013). For these reasons, as a first approximation, we propose that the number of distinct occupations necessary to maintain and reproduce a society is a reasonable measure of a society's complexity.

In our formal model (presented in the following sections) each occupation is considered a distinct machine. A society is considered a computational agent comprised of multiple interacting machines. Any human society is running numerous distinct algorithms simultaneously at each timestep. The reason why occupational specialization is appealing for our purposes is straightforward: an occupation can be considered an ensemble of algorithms for interacting with the environment and with other social agents.

Categorizing occupations, however, is no trivial undertaking. Several categorizations of occupational specializations already exist (Smith, Dooley and He). Each of these schema is organized hierarchically, with varying resolution (e.g. the U.N. scheme recognizes 436 occupations at the most detailed level, the E.U. scheme 3008); each also includes a description of the tasks or skills involved in each occupation. There is also a classification known as HISCO (Historical International Standard of Classification of Occupations), developed by researchers interested in the history of occupations (Leeuwen, 2022). Obviously, a simple count of occupations will be sensitive to the grain of resolution used to divide up occupations (e.g. engineer, chemical engineer, fertilizer synthesis engineer, and so on). But this is not cause for despair. First, it has already been shown by researchers interested in the scaling patterns of occupational diversity in modern cities that scaling patterns are independent of the degree of resolution (Bettencourt, Samaniego and Youn, 2014). Second, it would be possible to use word-embedding tools to identify the actual skills, functions, and tasks underlying occupational titles and to construct measures of the real distances in the networks of skills, functions, and tasks embodied in a given society's occupations.

We immediately see non-trivial ways in which this measurement might be too limited even on its own terms. First, as mentioned, not all occupations are created equal, and some occupations are presumably more computationally complex than others. A nuclear physicist may execute more algorithms, and more complex algorithms, than a beet farmer. It would be fruitful to explore this possibility by considering proxies for the computational complexity of a given occupation, such as years of training required. Second, two identical sets of occupations might interact in different ways, permitting meaningfully more or less complex computations. Therefore, the density and nature of interactions between occupations are important, as are the network structures that reflect how occupations communicate. Third, different algorithms, singly or in combination, may offer greater "bang per bit." In other words, some algorithms (or some occupations) may have a disproportionate thermodynamic effect on a society.

3.2 Specialization, scaling, and change. We observe two stylized facts that support the use of occupational specialization as a proxy. First, in the present, occupational specialization is strongly correlated with per capita wealth, underscoring the link between productivity and specialization (Figure 2).

Second, occupational specialization increases over time, mirroring other measures of social complexity such as largest city-size, levels of hierarchy, etc. We also believe that our proxy has the potential to capture recent (last 200 years) increases in social complexity in a more fine-grained way than existing measurements. Consider, as a first approximation, the number of unique words used by census respondents to describe their occupation between 1850 and 1940 in U.S. cities (Figure 3).

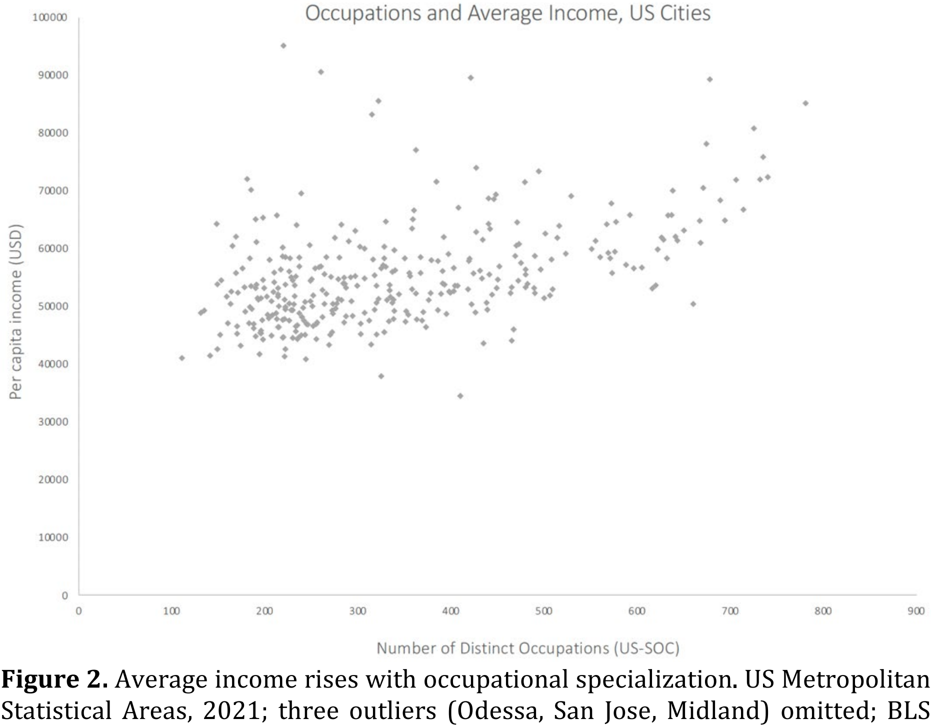

Figure 2. Average income rises with occupational specialization. US Metropolitan Statistical Areas, 2021; three outliers (Odessa, San Jose, Midland) omitted; BLS data, which uses the US SOC (Standard Occupational Classification) 2018 version, with 867 occupations at the most granular level.

Cities are "social reactors", (Bettencourt, 2013) creating dense networks for the flow of information. It has been observed that most socio-economic metrics of modern cities (from income to creativity to crime) scale according to a power law, whereas occupational diversity scales logarithmically, "illustrating the crucial role that diversity has in fostering a strong economy" (Bettencourt, 2014; Yang et al., 2022). This is an important insight, and indeed it suggests that occupational diversity is a critical ingredient in the growth of cities and the complexity of human societies more generally. In turn, occupational specialization is constrained by various factors, including the extent of the market and coordination costs (both of which are shaped by variables such as institutions, technology, and transportation infrastructure that are highly pertinent to understanding how a society processes information), in addition to what interests us most immediately, the stock of information itself (i.e. the stock of ideas or knowledge in the economist's terms).

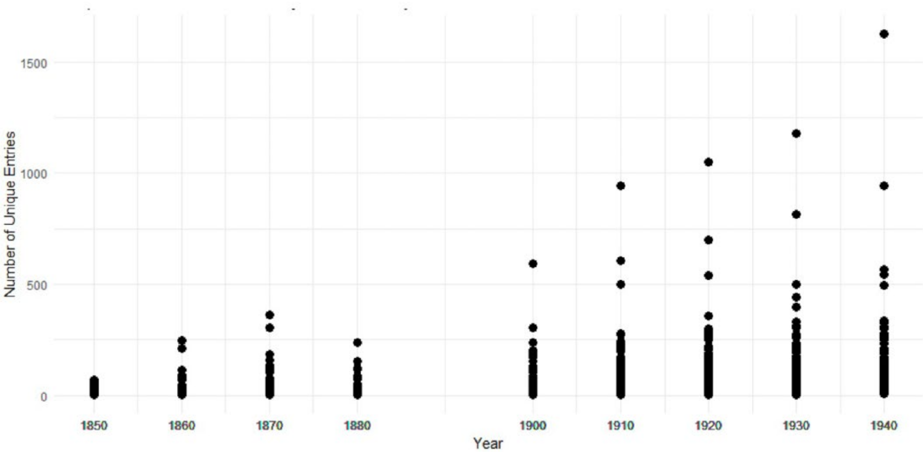

Figure 3. Number of unique words (minus articles, prepositions) used by census respondents in US cities to describe occupation; each dot is a US city. Data source: IPUMS (Integrated Public Use Microdata Series). This measure is independent of any classification system, and relies only on the respondent's direct answers, here in a one percent sample.

Consider the observation that occupational diversity scales logarithmically with urban population size (Figure 4, from IPUMS, data via the American Community Survey's Public Use Microdata Sample, a classification scheme that contains 530 unique occupations; see also [Bettencourt et al., 2014]).

Even a rough evaluation of the relationship between occupational diversity and urban population over time suggests interesting variation, however. Occupational diversity in US cities was distinctly lower in the 1940s (in the aftermath of the

Second Industrial Revolution) than at present, and lower still in the 1850s (in the midst of the First Industrial Revolution). In preindustrial times, specialization was even more constrained. It is not simply that urban populations were smaller in the past. For instance, if ancient Rome existed in today's urban network, with its population of 1M, we would predict it to have ca. 400 unique occupations, whereas in fact it has fewer than half of that (based on a count of occupational designations in inscriptions and texts from the Roman Empire; see also [Hanson, Ortman and Lobo, 2017; Kase, Hermankova and Sobotkova, 2022]). To some extent, these differences are sensitive to the difficulty of maintaining occupational classification at consistent resolution across time and space. But not entirely. Ancient Rome did not have chemical engineers, software developers, flight attendants, web designers, etc., and even at infinite resolution, real shifts have occurred because the stock of information has grown. We hypothesize (as in Figure 5) that the relationship between occupational specialization and population size has changed as a function of social evolution - that the computational-metabolic capacities of human societies have grown due to increases in the stock of information.

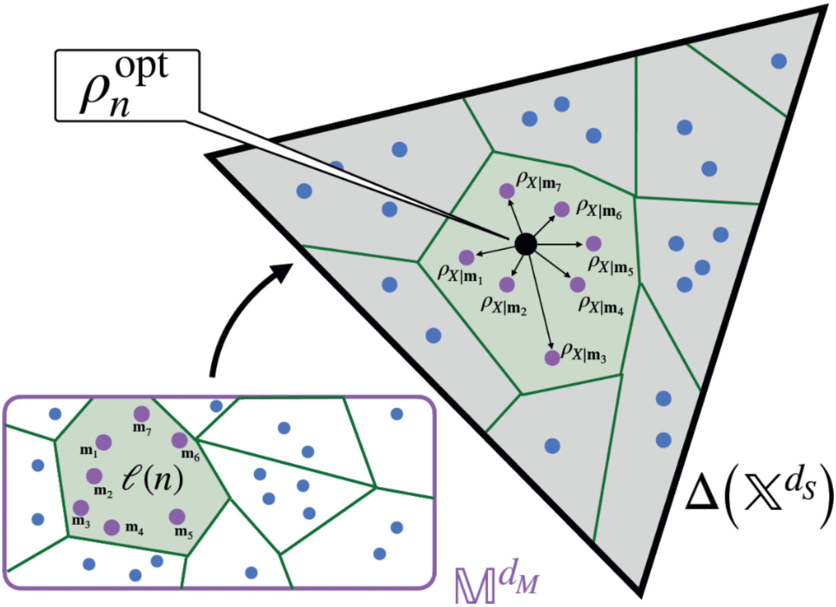

Figure 4. Urban populations (millions) and unique occupations, US cities, present (data source: IPUMS, based on the American Community Survey occupation classes).

In the future, we intend to build a dataset of occupational diversity in deep time in order to better understand the relationship between social evolution and the total stock of information. We plan to exploit past and present census data and to deploy text mining methods on large textual datasets (Wikipedia, Google Books, Hathi Trust) to analyze the history of occupational specialization.

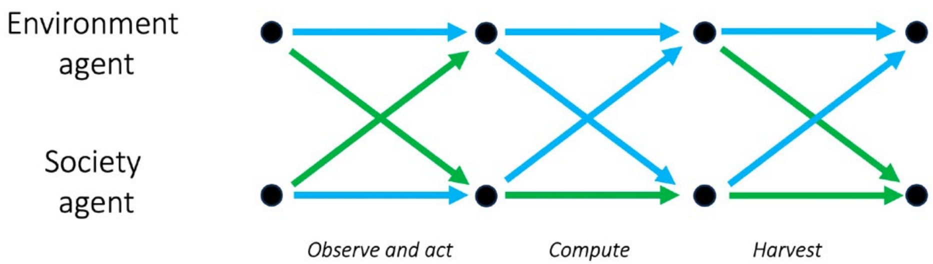

Figure 5. Schematic representation of social evolution and scaling of occupational diversity in cities by population (hypothesized).

Our hope is to be able to exploit this dataset to ask:

1) What factors account for increases in occupational diversity? Is it possible to distinguish the role played by market institutions, coordination costs, and the stock of information?

2) To what extent is there cross-cultural structure in the changing composition of occupational specialties over time? Could Gutman scaling be used to identify necessary or probable sequences of occupations in the evolution of complexity (in the manner of [Peregrine, Ember and Ember, 2004])? In other words, are societies universally likely to exhibit certain occupations in certain temporal sequences or other configurations?

3) Are there ways to exploit big textual datasets to understand the algorithmic complexity of specific occupations as well as the changing complexity of an occupation over time?

4) Is it possible to associate occupations with metafunctional categories such as production (e.g. metallurgist), consumption (e.g. entertainer), and information exchange (e.g. journalist), and to identify structures of change over time?

5) Is it possible to understand empirically how specific occupations interact in networks akin to the way that computers communicate (see following sections)?

6) Can changing occupational structures cast light on the macroevolutionary dynamics of human societies, revealing the relative and changing importance of processes such as convergent selection, conserved traits, horizontal transmission, contingency, the evolution of evolvability, etc. (as in [Enquist, Ghiralda and Eriksson, 2011])?

7) Using time series or cross-sectional analysis is it possible to understand the "bang per bit" (in the sense of [Kolchinsky and Wolpert, 2018]) of various occupations or algorithms or combinations thereof? In the long run, what kinds of increases in the stock of information have yielded the greatest changes in metabolism or complexity? Does the changing pace of change help illuminate the deep structure of social evolution?

We believe occupational specialization is a suitable first proxy because crude historical data are conceivably attainable for deep time and because occupations plausibly reflect algorithms. However, we hope a computational approach to social evolution could inspire complementary efforts. It would be valuable to understand qualitatively and quantitatively how information storage and processing technologies have changed over time; how deep changes in the structure of science and engineering alter the "bang per bit" of discovery; and how societies implicitly use algorithms to solve problems and how institutions shape social computation.

A computational model of societies

Changes (positive or negative) in the ability of a society to extract free energy from its environment ("harvest free energy") affect its abilities to run computations, which in turn affects how much free energy it can harvest, which in turn affects its abilities to run computations, and so on. (See the Bayes net illustrated in Fig. 6.) Empirical data show that this "compute-extract" loop is one of the primary drivers of the long-term dynamics of human social systems (Morrs, 2010; Shin et al., 2020). However, at present we have no theoretical understanding of this loop. Our ultimate goal — far beyond the scope of this preliminary paper — is to start to fill in this gap. It is important to distinguish our goal from the goals of already well-established bodies of work that also involve both CS theory and the social sciences. Our goal is to build a detailed model of the dynamics of a social system as it makes its decision. We wish to consider such models using the tools of computer science theory in general, and computational complexity in particular. In contrast, work in algorithmic game theory (Nisan and Ronen2001; Roughgarden, 2010) or on the computational complexity of computing Nash equilibria (Daskalakis et al., 2009; Papadimitriou, 2014) focuses exclusively on the best possible performance (typically on the worst-case problem) of a Turing machine at computing some abstraction of the end result of an infinite number of steps of a real-world social system computer. Our goal is to use CS theory to model the dynamics of a changing social system, whereas those other bodies of work instead use CS theory to derive results concerning the "static" properties of social systems, i.e., properties of their state at a single moment in time.

Our premise is that a potentially powerful way to investigate the computeextract loop is to model both the human society and the environment as computers, in the full CS theory sense of the term, that are iteratively interacting with each other. That might allow us couch the analysis in terms of the relationship between the attributes of two computational machines. It might also allow us to build upon the extremely rich body of theorems in CS theory. In the sequel, we will refer to each of those two computers as an "agent", where as needed, we clarify whether we are discussing just the Society agent or the Environment agent.

We draw inspiration for starting to pursue our goal from the field of artificial life. One of the central functions of living things is to store and transmit information, and to perform computation with that information, all the better to extract free energy from its environment. This conceptualization of biology led to the development of artificial life, first in theoretical work (Rogers, 1971; Sipper, 1998)[1] and then starting about half a century ago, in computational simulations (Aguilar et al., 2014; Langton, 1997). This has led to great insights into the fundamental nature of all living systems. Our goal is to try to kickstart a similar trajectory for archaeology and deep history, simply starting about 75 years after artificial life paved the way.

A central realization underpinning our approach is that there is major difference between communication and computation, even though social systems engage in both. Communication is all about trying to ferry bits from point A to point B with as little distortion as possible. It is about information transmission, and involves error-correcting codes, Shannon information, mutual information, channel capacity, and the other tools of the field of information theory. In contrast, computation involves information transformation, and is all about changing information, synthesizing different streams of information, etc. Rather than information theory, the mathematical field that is central to understanding computation is CS theory. While communication and computation are clearly closely related, and occur together in many natural systems (in particular in all social systems), they are fundamentally different. Unfortunately, this is completely unappreciated in the social science literature. Indeed, many people use the term "computation" and then proceed to invoke information theory. We view our introduction of some of the concepts of CS theory to the social science community as one of our most important (albeit trivial) contributions.

In the sections following this one we introduce a bare-bones model that we feel can serve as a seed for discussing more fully featured models. The hope is that this model might be simple enough to allow some formal, mathematical analysis, while at the same time have it be easy to enrich the model by adding insights from data sets and from domain expertise.

---

[1] It is sobering to appreciate that the artificial life community lauds Von Neumann for the creation of his "replicator", without any realization that this work of Von Neumann's was simply a restatement of Kleene's second recursion theorem of about a decade earlier. (Indeed, Kleene's recursion theorem is the foundation of the idea of a "computer virus" [Sipser, 1996].) Indeed, even though Kleene was a renowned logician, Von Neumann failed to give him any credit.

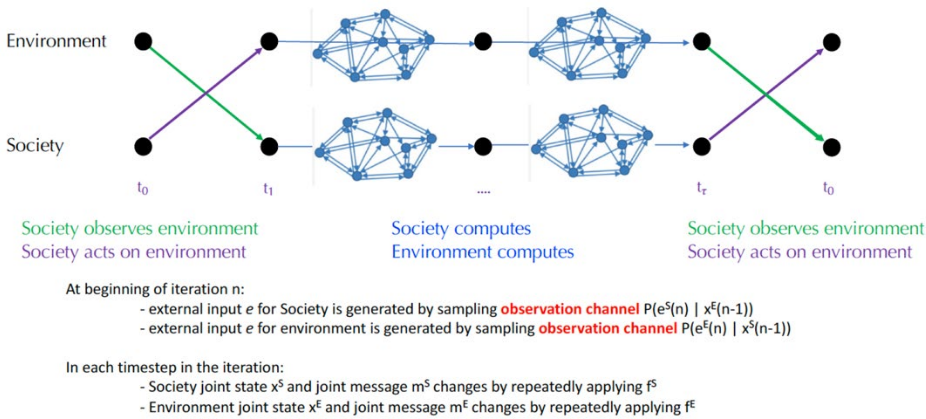

Figure 6. Bayes net model of a single iteration of human social computer interacting with an environment computer. Green arrows refer to actions actually taken by the Society Agent, as described informally in the captions.

The key feature of the computational process that we use is that each agent comprises a set of multiple interacting subsystems that collectively perform the computation within a single iteration. In particular, a human Society agent comprises (the members of) multiple interacting occupations, which collectively perform the computation the Society does to transform observations of the state of the Environment into actions back on that environment.

The next subsection briefly describes some of the scenarios considered in CS theory that were direct inspirations for our model for two agents that interact through an iterative process. In the following subsection we introduce our formal model of the dynamics of an agent within any single such iteration. The last subsection in this section then presents a set of many examples of our formal model that occur in the natural world. In the section after this one, we describe the the way that two agents jointly interact, in the fully general case. We also describe the physical meaning of the variables in this framework, paying special attention to how the framework captures features of the computation process of an agent.

In the following section we describe several ways to tailor our framework concerning two interacting agents to the special case of one agent representing Society and one the Environment.

Inspirations for our model of agent computation

The central feature of the model we introduce is that each agent (both the Society agent and the Environment agent) comprises a set of multiple *machines*. As described below, in this paper we view each machine in the Society agent as a separate occupation, which could easily be reconstructed to include technologies as distinct machines. (However, we believe the model extends far more generally, e.g., to usefully capture interactions among individuals, or among corporations, or among nations, rather than only among occupations and technologies.) We then have the "society computes" part of the interaction between the agents shown in Fig. 6 consist of some fixed number of "timesteps", in which all the machines in the Society agent can communicate with one another while also performing computations in each timestep based on what they currently know. In our model, that iterated sequence of interwoven communication and computation is how the agent as a whole computes. Accordingly, we call such an agent an instance of a "multiple communicating machines" (MCMs) model of computation.

There were several bodies of CS theory that were inspirations for the MCMs model. Perhaps the primary one is work in CS theory on communication complexity (Kushilevitz, 1997; Arpe et al., 2003; babai et al., 2003; Fischer et al., 2016). In standard communication complexity, there are k machines that only send single bits to one another at each timestep, and they must all end up calculating the desired function, $f(x_1, \ldots, x_k)$ [Rao and Yehudayoff, 2020]. There are several ways that the field of communication complexity has been extended to deal with more than two interacting machines though. One, called "simultaneous messages" (SM), assumes that (like in epistemic logic) each machine i can see x_{-i}, the joint input of all the other machines in the society (e.g., all the other bits being output by the environment). They cannot see *their own* input (Babai, 2003). Instead, the work on SM has every machine send a message to a *referee*, who sees none of the inputs directly, and that referee makes the evaluation of *f*. In our case, that sort of makes sense as a model of a hierarchy. In addition, there has more recently been work that is a "number-in-hand" variant of this original "number-on-forehead" model. This new work is called "private simultaneous messages" (PSM), and has every machine i receive *only* its own input x_i (Ball and Randolph, 2022; Fischer, 2016).

Communication complexity focuses on the minimal number of communication timesteps needed to jointly calculate a function. We're actually interested in something different; how well one can do (e.g., the probability of success) at calculating the function if one has a fixed amount of allowed number of timesteps. In addition, we are not interested in having every machine in the society agent calculate the exact same function.

Any modern computer continually runs multiple computations simultaneously, i.e., every modern computer is in fact a parallel computer. In this, they are exactly analogous to what happens in human social systems. Moreover, any such computation will often "fork" multiple new computations while it is still running, Any such fork increases the number of computations running simultaneously. This too is exactly analogous to what happens in human social systems.

In fact, how best to fork computations (or finish computations that are still running) is a central issue in the field of distributed computation. An important special case of distributed computation (albeit one without forking of new computations, etc.) is circuit complexity theory (Arora and Barak, 2009; Jukna et al., 2012; Spiser, 1996), which is closely related to our problem. Perhaps the easiest way to understand what this body of work is about is to consider Boolean functions. A "Boolean function" f over n bits maps any possible sequence of n "input" bits to a single "output" bit. A "Boolean circuit" is a way to implement such a Boolean function.

To do this, a Boolean circuit arranges "gates" chosen from some fixed set of a few types of logical gate, like the logical gates AND, OR and NOT, and feeds those into one another, in an acyclic fashion, to implement the Boolean function. Concretely, a "Boolean circuit" is a set of such gates that are fed into one another, iteratively, through successive layers. The first layer is the sequence of n bits that are the inputs to f, and the output bit is given by the output value of the (single) gate in the last layer of the circuit.

There are many measures of the circuit complexity of any given function f. Examples include the minimal number of gates of any circuit that can implement that function for all possible input bit strings, and the minimal number of layers of any circuit that implements that function for all possible input bit strings.

Other work in circuit complexity theory considers circuits where the individual gates perform far more (or less!) rich information processing than simple Boolean primitives like AND, OR and NOT. The MCM setup can be seen as such a circuit, where the individual gates are finite state automata, where each gate at each layer has a specially identified parent gate in the previous layer, and sets its initial state to be the ending state of its parent gate before processing all the inputs it receives from the other gates in that previous layer. Finally, it is worth pointing out the work in (Hastad et al., 2017), which considers circuits that are provided with a particular computational task (related to "Sipser functions"). For that task, if we measure performance by considering the average case over all possible input strings, it turns out that reducing the reducing the number of layers in the circuit by 1 can drastically decrease performance.

MCMs can also be viewed as a special type of a multi-agent system (Romanowska, Wren and Crabtree, 2021). In addition, the MCMs model can be viewed as a special case of a collective (Wolpert, Lawson, 2002; Wolpert et al., 2023; Tumer and Wolpert, 2000; 2002), where rather than an (exact) potential game (Monderer and Shapley, 1996; Candogan, 2011), we have a team game (Marshak and Radner, 1972), and the precise form of the communication among the members of the collective must obey multiple constraints. In addition, cellular automata (CA [Wolfram, 1984]) can be viewed as a special case of MCMs. This is done by identifying each cell in a CA as a separate machine in an associated MCM. In general, an MCM has a far richer graph of inter-cellular direct communication, and a larger state space of each cell, than any CA. In fact, so impoverished is the computational ability of cell in a CA, that information storage and processing is typically identified only with inter-cellular dynamics, not with behavior occurring inside any individual cells (Lizier et al. 2014). This is in stark contrast with the case of MCMs (and human societies).

A model that is more directly related to MCMs compared to CAs is known as the "CONGEST model" of distributed computation (Hirvonen and Suomela, 2021). The CONGEST model can be viewed as a type of MCM in which every machine can send log(n) bits of information to each of its neighbors in the message graph, in each timestep, where (for example) n might be the number of machines in the overall system. However, in contrast to MCMs, in typical CONGEST models there are no restrictions imposed on the computational power of the individual machines. Also, the overall system's performance (its "complexity" to use the CS theoretic term) is measured by the number of timesteps the overall system requires to solve the provided problem perfectly. In contrast, as elaborated below, with MCMs performance is measured by the quality of the overall system's response to a provided problem after a fixed number of timesteps[2]). Finally, the analyses of the CONGEST model usually consider worst-case over inputs to the overall system, whereas with MCMs we are almost always most concerned with the average performance over the possible inputs.[3]

A more direct inspiration behind the MCMs model than CONGEST was the idea of formalizing the computation done by a human society during each iteration as a computation by a Turing machine (TM [Arora and Barak, 2009; Li and P., 2008; Sipser, 1996]. Each iteration would start with a new input provided to that TM by the environment, and that TM would then evolve in the usual way through a succession of timesteps. After it completes the next iteration would begin. Particularly compelling is to consider a TM with multiple work tapes, where each tape models a different occupation or technology in a society.

One difficulty with this approach is that TMs *halt*. That means that there is no way for a society modeled with a TM to have memory from one iteration to the

---

[2] In some ways this performance measure of MCMs is similar to what is called "approximation complexity" in CS theory (Arora and Barak, 2009). See also (Hastad et al., 2017).

[3] In some ways this feature of how we measure performance with MCMs is similar to what is called "average-case complexity" in CS theory (Arora and Barak, 2009).

next. A natural way around this shortcoming is to suppose that the input to a society TM at the beginning of each iteration includes its own state at the end of the previous iteration, in addition to getting information about the state of the environment. However, unless that "state at the end of the previous iteration" is extended to include the ending states of the work tape, there would be loss of memory of those ending states from one iteration to the next. So for example each occupation would forget what it had previously been thinking. This problem can be circumvented if we instead model society as a monotone prefix TM (MTM), i.e., a TM that never halts, iterating forever, where we stipulate that each iteration lasts for some fixed number τ of timesteps before the next iteration begins — and with it arrives the next input from the environment.[4]

However, even once we replace TMs with MTMs that run for exactly τ timesteps in each iteration, there are still two other substantial restrictions that we need to address if we want to use our formalism to model the computation by a human society. The first is that there are restrictions in how much the heads on each work tape can move along that tape in a single timestep in an MTM, whereas there is no analogous restriction on how much the state of an occupation or technology can change in a single timestep. The second restriction is there are no direct ways in an MTM model to allow different work tapes to talk directly with one another.

The model of the society agent that we investigate in this paper is a fully formal version of an MTM that runs for exactly τ timesteps in each iteration, with the two restrictions mentioned just above removed. We call this model a "multiple communicating machines" (MCM) model, using the term "machine" rather than "work tape", since the machines do not have restrictions on how they can communicate with one another.

Another important aspect of our investigation is to determine how a society reacts to information from its environment to best act upon that environment. Our modeling of this aspect of human social systems draws inspiration from the field of feed-back control. Feed-back controllers are artificial systems that react to stimuli in their environment. They underlie all of modern civilization's functions; without them our modern society would immediately grind to a halt. (There are many feedback controllers in the reader's home, in fact.) Moreover, aside from very simple feedback controllers, like thermostats, all feed-back controllers use embedded artificial computers to decide how to analyze their stimuli from their environment, and then act back on their environment. Again, the analogy between computers and human social systems is exact.

Ultimately, we would like our MCM model to also serve as our model of the computation done by the environment. The motivation for using the same model

[4] Note that imposing such a value τ is similar is similar to how an upper bound on the number of allowed timesteps is imposed in the definitions of the usual computational complexity classes, e.g., P or NP

for both systems is that it might allow us to directly compare the computational characteristics of the society agent and the environment agent. We might then be able to analyze how their relative computational powers affects the ability of the society agent to grow through time by interacting with the environment agent. In the sequel though, we will focus on using MCMs to model human societies.

Formal definition of Multiple Communicating Machines model

In this subsection we first formally define MCMs, and then in the following subsection we provide a few examples, illustrating the role the definitions will play in our investigation of the interaction between human societies and the environment. Of course, in specific cases it might be worthwhile to investigate formal models where some of the following details are simplified (or even removed), and / or some of the following details are elaborated in more detail. We will refer to such models as **simplified** MCMs.

1. There are N **machines** (or "sub-agents") in an MCM, indexed by values $v \in V$. The state space of machine v is X_v, with elements x_v. The joint state of all the machines is written as $x \in X$.
   We will abuse terminology and refer to "cardinality of X_v", or write "$|X_v|$", even though we mean different things depending on the precise type of space X_v. For example, if X_v is finite, then $|X_v|$ means the number of elements of X_v, while if it is a Euclidean space R^n it means the dimension n, and if in fact the Cantor cardinality of X_v is at least $\aleph_2$, then $|X_v|$ means the Cantor cardinality of that set.
2. There is a directed graph G with a set of N nodes, V, each of which is identified with one of the machines. An edge from node v to node v' is written as (v, v'). The parents of node v are written as pa(v), and the children are written as ch(v). G is called the **message graph** for reasons that will become clear.
3. There is a set of possible **external inputs**, $\mathcal{E}$, with elements e, which for simplicity is the same for all machines v.
4. There is a set of possible (machine) messages, M, with elements m, a set which for simplicity is the same for all machines v.
5. As shorthand, for each machine v, a vector of messages $m_{v'}$ indexed by machines $v' \in$ pa(v) is written as $in^v \in IN_v$. (Intuitively, during each timestep, machine v will receive a message vector in^v from all the machines who are the parents of v in the message graph.)
6. There is a machine-indexed update function

$$f_v : X_v \times \mathcal{E} \times IN_v \rightarrow X_v \times M_v \tag{1}$$

The update function f_v has several arguments, which we consider in turn.
The first argument of the update function is just the machine's own state.
The second argument will be identified with an external input e, which concerns the other agent (the one not containing the machine v), and which machine v sees during each timestep in an iteration. Note that this external input to v is an extra input to v, in addition to the messages from some of the other machines in the agent containing machine v. As described below, within any given iteration of the two-agent Bayes net, the external input to agent v will not change from one timestep to the next.
The last argument is the vector of all the messages that v receives from other machines in v's agent. In general, these will change from one timestep in an iteration to the next.
The first output of the update function of machine v is a state of machine v.
The second output of the update function of machine v is a message that machine v sends to some other machines (as specified in the message graph).
For simplicity we assume that each f_v is a total (computable) function, guaranteed to finish between one timestep and the next, for any triple of arguments to that function.

7. To perform its computation, the MCM runs for a total of τ **timesteps**. In each timestep t, all machine in the MCM runs their associated update functions simultaneously. The state of each such machine v represents whether it has been activated (and potentially, how active it is, depending on our precise choice of X_v). Machine v uses the associated component of the output of its update function to set its state in the next timestep, $t + 1$. The other component of the output of v's update function is a single message, m_v. This message forms the associated input argument in the next timestep to every machine $v' \in \mathrm{ch}(v)$.

Note that for simplicity, we suppose that the messages that each machine v sends to all the machines in ch(v) each time it runs its update function (i.e., in each timestep) are identical, not varying among the receiving machines in ch(v). On the other hand, this single identical message that v sends can change from one timestep to the next. And of course, it will vary among the machines v in general.

Examples of MCMs

A concrete, physical example of a naturally occurring system that we can model as an MCM is a genetic regulatory network (GRN). In this example, each machine in the MCM model of the system is identified with one of the genes in the network. The messages in the MCM would be identified with transcription factors the genes send to one another. So for example, if transcription factors are transmitted via diffusion, then the directed edges in the message graph would be defined as the set of all pairs of genes (v_1, v_2) where v2 responds to the transcription factor released by v_1. (Note that in general this will be a proper subset of the edges in a fully connected graph among the same set of machines, since not all genes respond to all other genes.)

The external input e could then be some vector of counts of certain biomolecules from outside of the cell. In real-world examples of GRNs, often there is no natural way to specify some fixed number of timesteps, τ. This illustrates how in a real-world scenario, we might want to consider a simplified MCM rather than a full one (cf. Section 4.2).

Another example of a system in the natural world that can be modeled with an MCM is a hunter-gatherer band. In this example, each machine in the MCM model of the system is identified with an individual person in the band. The messages in the MCM would be identified with things that the individuals say to one another (or in some other way signal to one another). The external input e could then be information that the members of the band collectively gather about their environment. In contrast to the example involving GRNs, quite often there will be a natural choice of τ in the case of a hunter gatherer band, e.g., if the band is hunting and must complete the hunt before the sun goes down.

A related example — central to this paper — is a (complex) human society. There are many different versions of this example, depending on what societybased informational systems we identify with the separate machines. The example of such a process that we focus on in this paper is a human occupation. There are many other processes that we include in our model as well though. For example, we could identify various technologies with some of the machines — in this scenario, the “messages” back and forth between individual occupations and technologies would model the uses of that technology by members of that occupation. Another important process is information storage systems, e.g., libraries. So for example we could use messages exchanged between individual occupations and libraries as models of members of those occupations storing and / or retrieving information from those libraries. Other obvious examples of individual machines in a model of a human society include specific bureaucracies, specific corporations, etc.

We can also model an environment that a human society interacts with as an MCM. The machines in that MCM could be systems like the weather, plagues, plants that might be receptive to domestication as food crops, etc. (As tongue-incheek, crude intuition, it may make sense to identify each machine in the environment MCM with a phenomenon in the natural world that polytheistic religions ascribed a deity to.) The external input of the Society MCM would be the observation it makes of the state of the Environment MCM.

In this example of a society and environment, we might also invoke the ChurchTuring thesis (Piccinini and Maley, 2010; Piccinini, 2011) and require that f_v be Turing-computable for all v. Also on physical grounds, no matter what the computational power of f_v, it makes sense to require that the time to run it is very small on the timescale of the MCM's timesteps, for any input f_v might receive.

Yet another physical example is a spin glass (or more generally, Potts model, or some such) undergoing Glauber dynamics. In this example we identify each machine in the MCM with a separate spin, with the message graph reflecting the coupling terms in the Hamiltonian of the overall MCM. The external input could then be a magnetic field.

Moving on to the digital world, note that MCMs can be viewed as instances of parallel computers, which are ubiquitous in modern-day digital computers. To see this we simply identify each machine in the MCM with a different processor in a parallel computer, with the messages among those machines identified with interprocessor communications. (We could augment this with details about shared memory among the processors, rules for how access conflicts are managed, etc.) In the same vein, MCMs likely have connections with circuit complexity theory, which is closely related to the theory of parallel computation (Arora and Barak, 2009; Sipser, 1996).

As another more abstract example of an MCM, suppose that we allow τ to be infinite, and have e be an arbitrarily long sequence of bits. Suppose as well that every machine v has an N-valued counter n_v in it (i.e., n_v is a component of the vector x_v), and that n_v increases by 1 in every timestep for every machine. Finally, assume that for all machines v, the associated update function f_v can only depend on bit number n_v in e. In this case, each machine in the MCM is a Mealy machine. So the full MCM is a set of N simultaneously communicating Mealy machines.

Finally, it's worth pointing out that all of the motivating examples given in Section 4.1 (which do not occur in the real world) can also be formulated as MCMs. To do this though the machines in the MCM must have countably infinite state spaces (uncountably infinite, in the case of a CA) and in the case of communication complexity, the machines would typically be allowed super-Turing computational power.

Taken together, these examples demonstrate that while there is a fair amount of mathematical structure in the definition of an MCM, it can still model a very wide range of real world distributed computational systems. In addition, little of that mathematical structure is superfluous, in that most of it plays an important role in the examples; simplified MCMs are often too impoverished to capture important features in these examples.

Nonetheless, we emphasize again that many of the *other* features of these examples would be mostly unaffected by associated simplifications of the MCM model. If some such features are the primary interest of the modeler, then it makes sense to simplify the MCM model accordingly.

Interacting MCMs

In actual physical systems, the exact same mathematical structure describes two processes which are often misunderstood to be intrinsically different. The "ob servation" of agent B by agent A is an instance of the state of agent B at time t affecting the state of agent A at some later time t' . (Formally, changes to B_t result in changes in the conditional distribution $P(A_{t'}|B_t)$.) Similarly, an "action" by agent B on agent A is an instance of the state of agent B at time t affecting the state of agent A at some later time t'. So it is an "observation" by agent A of the earlier state of agent B. We can illustrate this with the example of a human Society MCM interacting with an Environment MCM. In this case, the external input of the Environment MCM would in part reflect the action that the Society MCM takes on the Environment MCM. (Such actions can range from types of niche construction like constructing dwellings in the environment, to sowing seeds in the environment for future harvest.)

In this section we present a general framework of an "iteration" of two interacting MCMs. At the beginning of such an iteration each MCM observes the other one (or equivalently, each acts on the other one). They then each carry out an independent computation, by repeating their respective update functions (as described in Section 4). This gives us a basic, pared-down model of interacting MCMs, with an observation process that also models control and action of one MCM by another.

The following section, Appendix A, presents a model built on this basic framework of interacting MCMs, tailored to the special case of one agent representing a complex human Society and one representing the Environment. This model adds extra structure that formalizes how the joint state of those two MCMs at the end of an iteration can determine the free energy harvest by the Society from resources in the Environment, with the resultant modification of the parameters of the Society MCM.

See Section 6.3 for some other ways to extend this basic model of interacting MCMs.

The joint interaction between MCMs

Each of our two MCMs is defined by an associated collection,

$$\left\{N, G, \tau, \{M_v\}, \mathcal{E}, \{f_v\}\right\} \tag{2}$$

When we need to consider the two agents simultaneously, we will index these features accordingly, e.g., writing τ^A for the value of τ for agent A.

Loosely speaking, an entire run of τ timesteps of each MCM (one MCM modeling the society and one modeling the environment) is interpreted as the computation by that MCM occurring in each iteration of the Bayes net depicted in Fig. 6. Those computations of the two agents are done in parallel. So for example, if $\tau^A = \tau^B = 1$, then the two agents interact with one another continually, i.e., interact with one another just as frequently as the individual machines in each agent interact with one another. More generally though, τ^A and / or τ^B may exceed 1, in which case the associated agent(s) spend some timesteps computing by themselves, without interacting with the other agent, before they interact again with that other agent.

The external input to each machine v in an agent is interpreted as the result of the observation of the other agent that the machine v makes at the beginning of such an iteration. This is why the external input does not change from one timestep to the next in the formal definition of an MCM in Section 4.2. As an example, this external input to some machine v in the Society agent is the attributes of the environment observed by the members of the machine v at the beginning of an iteration. Note that because of this feature of the interactions, if the values τ^i of both agents i are greater than 1, then the agents are performing their respective computations faster than their interactions with the other agent.

We now present the details of our model of the interaction of two MCMs :

1. An **iteration** of two interacting agents (MCMs), A and B consists of a prefixed set of τ^A timesteps for agent A and τ^B timesteps for agent B. We adopt the convention that time in each iteration starts at 0, followed by timestep 1, timestep 2 (if the iteration hasn't ended by then), etc.

   The Bayes net in Fig. 6 is a stylized depiction of an iteration, omitting the computation of the environment agent, and omitting details of the computation by the society agent.

There are a countably infinite number of such iterations succeeding one another, forming a sequence. We adopt the convention that the joint state of the machines in the two agents after the last timestep of iteration $t - 1$ is the same as their joint state before the first timestep in iteration t.

2. In the first timestep of iteration t, agent A receives its external input from the final (timestep τ^B) joint state of the machines in agent B at the end of iteration $t - 1$, according to some conditional probability distribution,

$$P(e^A(t) \mid x^B(t - 1)) \tag{3}$$

   This is the **observation channel** [20] of the state of agent B by agent A. This observation takes place in the first leg of Fig. 6. (See also Fig. 7.)

   Note that the external input to (all of the machines within) agent A will be the same value for all the timesteps in a given iteration t. This contrasts with the messages, which may change at each timestep within that iteration. This difference reflects the fact that messages occur as part of the single computation that an agent A does in response to the observations made by the machines v in agent A concerning the other agent, and these observations by the machines in A are identified with the external inputs to those machines.

   As an example, consider the case of a Society agent A and Environment agent B. In this situation, $P(e^A(t) \mid x^B(t - 1))$ models the observation that the members of the occupations in the Society agent are making throughout iteration t of the state that the Environment had at the end of the previous iteration, $t - 1$. (Note that notationally, in this conditional distribution the arguments of the x's specifies the iteration, not the timestep within the iteration.) In the case of a environment, agent B, $P(e^B(t) \mid x^A(t - 1))$ models the effect throughout iteration t on the Environment of the joint action that the members of the Society agent took on the Environment at the end of the preceding iteration, $t - 1$.

3. After this first timestep, all machines in A repeat their respective update functions for the total of $\tau^A{-}1$ remaining timesteps, at which point iteration t has ended.

Mutatis mutandi for the external input to agent B from the final state of agent A at the end of the previous iteration, and for its ensuing dynamics.

Note that if $\tau^A \neq \tau^B$, then we implicitly assume the two MCMs compute at different speeds, so that their iterations both end at the same (wall clock) time, at which they take each take their actions / perform their observations. For pedagogical simplicity though, in much of the discussion below we will take if $\tau^A \neq \tau^B := \tau$.

4. There is also a starting distribution over the joint state of all the machines in both agents, which gets propagated in the usual way by the deterministic update functions of the two agents. For simplicity we have not written it explicitly.
5. To simplify the exposition in Appendix A below, we will implicitly assume that the observation channel of each of the two agents is a symmetric noisy channel. So each of those channels is parameterized by a single real number, setting the noise level. We write those two noise levels as σ^A and σ^B, giving the noise in the observation channel of agent A and of agent B, respectively.

We will refer to the set of quantities concerning agent A that are specified in Eq. (2), together with the noise level σ^A, as the **computation parameters** of that agent. As described below in Appendix A, all but one of the computation parameters of the environment agent do not change from one iteration to the next.

Figure 7: Expanded view of the Bayes net model in Fig. 6 representing an iteration of a human social computer interacting with an environment computer. Green arrows refer to actions actually taken by the Society Agent, and purple arrows refer to those actually taken by the environment as described informally in the captions. The messaging back and forth by the machines in each agent are informally represented by the small directed graphs in the middle leg.

However, as also described in Appendix A, the computation parameters of the society agent *do* change between iterations, as specified by the "evolution function" of that agent.

Physical significance of an MCM's computation parameters

The computation parameters of an MCM determines its maximal computational power. In particular, note that everything else being equal, the greater N^A is, the greater the computational power of agent A, (just like increasing the number of processing units in a parallel computer increases the power of that computer).

Note also that the bigger the state spaces of the machines in an MCM is, the greater the memory space of each of those machines, and so the greater the computational power of the MCM.

Similarly, the greater $|M^A|$ is, the more information the machines in the MCM can send one another in every timestep. Accordingly, increasing $|M^A|$ increases the computational power of that MCM.

Along the same lines, the maximal fan-out of the message graph of the machines in the society agent will affect how much computation it can do in each iteration. As an example, several of the "complexity characteristics" recorded in the Seshat data set of the evolution of ancient societies (Turchin et al., 2018) were identified in (Shin et al., 2020) as capturing the "computational power" of those societies. These complexity characteristics included the sophistication of the writing system of the society, the sophistication of its monetary system, and the sophistication of its transportation network. The direct analogs of these in the interacting MCMs model is the parameters $|M^A|$ and the fan-out of the message graph of agent A.

In addition, in the case of two interacting MCMs, the smaller σ^A is, the more accurately agent A can ascertain the starting state of agent B at the beginning of an iteration. Similarly, the smaller σ^B is, the more accurately the agent A can set a control action that it takes on agent B.

In this regard, note that the action that agent A takes on agent B at the beginning of iteration t is based on the computation that agent A completed in the *previous* iteration, $t - 1$. That computation in turn reflects the observation that agent A made of the state of agent B at the beginning of iteration $t - 1$. So loosely speaking, that agent A's computation is how it produces its answer to the following question: "If you observe {blah} about the state of agent B at a certain time, what is the action you should take to optimally affect the state that agent B will have by the time you finish answering this question?" (The precise meaning of "optimal" in the context of a society-environment pair of co-evolving agents is discussed below.)

In real human societies, the population size is an extremely important variable. Indeed, the poorly chosen term "scalar stress" (which is almost synonymous with population size) is ubiquitous in numerical data analysis in archaeology. Indeed, variables involving population size are rife in the 50 or so data fields recorded in the Seshat dataset (Shin et al., 2020; Turchin et al., 2018). Population size does not occur explicitly as a variable in the definition of an MCM. However, it does arise implicitly, in several variables that are in that definition.

As an example, everything else being equal, if the population of "laborers" in a human society who are acting directly on the environment increases, then the aggregate actions of those laborers on the environment is more precisely chosen, with less variability. That can be interpreted to mean that raising the overall population of the laborers in a society results in shrinking σ^E.

Similarly, if the number of members of a specific occupation in a human society increases, and if those members are at least somewhat distinguishable from one another, then that means that the joint state of the members of that occupation increases. In terms of the MCM formalism, that means a growth in the state space X_v for the machine v that represents that occupation. Along the same lines, suppose the number of members of a specific occupation increases, and they need not all send or receive messages to the exact same set of other machines. In terms of the MCM model, this would mean that increasing the population of that occupation causes a rise in the fan-out (out-degree) and fan-in (in-degree) of the node in the message graph that represents that occupation.

As a final comment, note that in the bare-bones version of the MCM framework presented above, the external input e of each machine in an agent does not change from one timestep to the next during an iteration of the agent. This reflects a "separation of timescales" between the processes of computation within each agent and the processes of each agent observing the other agent.

Features of our co-evolving MCMs model and ways to enrich this model

Sections 4 and 5 present a high-level framework for describing interacting agents, with some discussion of the possible physical meaning of the variables arising in that interaction. In many ways this is a minimal framework, providing no more than a way to capture the features discussed in Section 4.1. As described at the end of Section 2.2, one of our goals in formulating this framework was that it not have too much structure to be amenable to mathematical analysis, and consequently be "coarse-grained" when considered as a model of an actual physical system.

However, one of our primary goals is to use this minimal framework to investigate the MET of the Anthropocene (and more generally of human societies in the latter part of the Holocene). To do that we need to introduce extra mathematical structure, specifying how the parameters and functions defining the Society agent evolve in time, in response to the society's interactions with the Environment agent. As discussed at the end of Section 2.2, this extra structure results in a "moderate-grained" model of the interactions between complex human societies and their environment.

There are very many details that seem unavoidable in such a moderate-grained model, unfortunately. To keep the discussion in the main text coherent, we present one such moderate-gained model in Appendix A, rather than here in the main text. For simplicity and clarity, from now now on we no longer talk generically about two co-evolving agents A, B. Instead we consider the specific pair of a Society and Environment co-evolving agent, which are the focus of our study. We will distinguish the computation parameters and functions of those two agents with the labels *S* and *E*, respectively.

The underlying idea of our moderate-grained model is that it takes energy to maintain high levels of each of the computation parameters of the Society agent, i.e., there is an energy-cost function associated with each computation parameter. In addition, at the beginning of each iteration the Society MCM will "harvest" a certain amount of energy from the Environment MCM. We refer to this as the **gross free energy harvest rate** (GFER). The GFER must cover all the thermodynamic costs of the society's computation (as well as many other thermodynamic costs, e.g., maintaining homeostasis of the members of the society).

The coupling between the thermodynamic harvesting and the thermodynamic expenditures is achieved via the computational power of the society. The GFER is determined by how well the Society can control the Environment MCM, acting on the Environment in a way that causes it to have a desired state in the future. In turn, the maximal ability of the Society agent to exert that control is determined by the best computation it can do, and so by the values of the computation parameters of the Society MCM. In sum, how much computation the Society can perform and how well it can use the results of its computation to control the Environment determines how much energy the Society can harvest from the Environment. That energy harvest in turn determines how much computation the Society MCM can perform, and how well it can can use the results of its computation to control the Environment.

One high-level point to keep in mind is that the MCMs model is a *model*. The precise mapping of the variables in the model and their dynamics to associated phenomena in the real world is not specified, and may in fact vary from one use of the model to another. In particular, the "energy harvest" process at the end / beginning of an iteration takes one timestep in the model. However, in the real world, it may take an extremely long time, e.g., if a semi-static thermodynamic process is used to maximize efficiency. All such details are elided below.

There are numerous other salient features of the co-evolving MCMs model which are worth elaborating, in addition to those that go into our "moderategrained" model. There are also many features that we may want to introduce to extend this model. In the rest of this section we present both those features of the MCM model.

Features of the interacting MCMs model

There are several features of the interacting MCMs model that are worth emphasizing.

1. If τS is large while τ^E is small, then the Society does a lot of computation between two successive times when it acts on the Environment — and the Environment does *not* do much computation between those times. So in this situation, the Society agent is computing what to do quickly on the timescale of the dynamics of the Environment.

2. The fact that our model allows the technologies to change from one iteration to the next might provide a way of incorporating Guttman scaling (Forrester, 2009; Peregrine, 2018; Peregrine et al., 2004) into our analysis, i.e., of incorporating the idea that some sets of new technologies can only be added in a particular sequence, with each step in that sequence enabled by an increase in the GFER at the beginning of an iteration. One way to integrate this feature into the model is to ensure that the dynamic functions evolution operator only adds machines with certain capabilities in such a Guttman sequence.

3. One of our primary goals in formulating the interacting MCMs model is to use it to investigate the explosive escape from Malthusian traps that has happened in human society in the last half century, starting in the West. (Some authors refer specifically to the gap that opened up between the West and China as "the Great Divergence" (Pomeranz, 2000), whose dating is highly disputed [Sheidel, 2019]). In terms of our interacting MCMs model, that escape corresponds to a super-linear rate of growth of GFER.

   Accordingly, to investigate that escape we need the pair of the evolution function ϕ and harvest function ψ to have the property that — for a certain very restricted set of their input values (i.e., a very restricted set of joint values of the GFER and the computation parameters of the society agent) — the effect of the output of ϕ on the joint dynamics of the Society and Environment agents causes yet higher GFER in the next iteration. In particular, it would be very interesting if in the very next iteration, the GFER increases faster than population size, so that the GFER per capita rises, before the population increases enough in the subsequent iterations without any more rise in the GFER so that the GFER per capita falls back to its original value. This would be a formalization of escaping a Malthusian trap

4. Ideally, we want there to even be situations, i.e., joint states of the two agents and associated GFER, such that once the pairs of agents hits one of

those joint states, ϕ together with ψ causes a "run-away" dynamics, where the GFER starts increasing exponentially faster than the population size from one iteration to the next for a sequence of multiple iterations.

Phrased differently, we want our model, and in particular this pair of functions, (ϕ, ψ), to endogenously capture a phenomenon whereby once a threshold is reached, one has a "punctuated equilibrium", of a sort, i.e., a discontinuous leap to a new level of computational power. In our case, to model the current major evolutionary transition, we want these leaps to cascade, i.e., we want each one quickly lead to a next one, etc., in a quickening avalanche of such leaps.

5. Note that this property may already be built into our model. As the Society agent adds new machines from one iteration to the next, the associated message graph may undergo phase transitions, from having only small, disconnected components to having giant components, and eventually to being fully connected. This would correspond to a "jump" in the computational power of the Society MCM. Similarly, as τ^S increases and / or the number of machines in the Society MCM increases and / or its message graph's fan-out grows, the computational power of the Society MCM increases.

   For a fixed Environment MCM, such increase in the computational power of the Society MCM might go through a "phase transition", suddenly allowing it to predict the future state of the Environment MCM far more accurately than before. Alternatively, reductions in σ^E might cause a jump in how well the Society MCM can control the future state of the Environment MCM, thereby simplifying the computational difficulty of predicting the future state of the environment.

   In all of these cases, the jump in the ability of the Society might lead to a jump in the value of the GFER that the Society agent acquires at the end of the iteration. (For example, one would expect such a jump in the case of the mutual information harvest function, described in Appendix A.1.) These jumps might then "feed on themselves", resulting in an accelerating sequence of jumps in the value of the GFER, and therefore in the energy usage per capita of the society.

6. Note that it is pretty much incontrovertible that societies don't really consider the costs associated with the future, only the benefits. Moreover,

they typically only look a very short distance into the future. So rather than than try to increase something like the expected future-discounted GFER (which depends on future costs), they try to increase the immediate (next iteration) GFER. For us this might be an advantage, since it presumably substantially simplifies the approach to setting the computation parameters and dynamics functions and allocation distribution to be optimal.

7. More generally, independent of what human societies actual do or don't try to optimize, there might be some rich mathematical analysis even if we restrict attention to a single iteration of the co-evolving MCMs, without considering the evolution function of the society agent or the nature of the environment agent's dynamics across iteration boundaries.

   In particular, if the environment were dynamically simple — that is, if it were simple when viewed as an MCM — then intuitively, one would expect that there would be no benefit to human society of increasing its computational power over multiple iterations. It would be interesting to either establish that this is the case, or show that it is not.

Extending the MCM model

The definitions in Section 4.2 only provides the most basic type of MCM. Many of the features of this basic type of MCM are chosen for pedagogical simplicity, and to open the possibility of mathematical analysis (e.g., of scaling behavior). This definition can be enriched in many ways, to define a more complete (and therefore more complicated) type of MCM.

1. In the basic type of MCM, the messages that v sends to all the machines in ch(v) each time it runs its update function (i.e., in each timestep) are identical, not varying from one receiving machine to the next. In the full model, that restriction is relaxed, allowing each machine v to send a different message to each of its child machines in the message graph.

2. Similarly, the message spaces of the machines in an MCM are identical in the basic type of MCM, whereas they are allowed to differ from one another in the more complete definition of an MCM.

3. In the basic type of MCM. the messages that any machine v sends to all the machines in ch(v) each time it runs its update function (i.e., in each timestep) are identical. This means that we do not allow a machine to

"tailor" its messages to the recipient. A natural way to enrich that simple MCM model would be to allow such tailoring, i.e., allow v to send different messages to the machines in ch(v) every time it runs f_v .

4. In the real world there will not be an update (deterministic) *function*, f_v, but rather an update *conditional distribution*. Again for pedagogical simplicity, the simple MCM model does not have this flexibility, i.e., it presumes that all such distributions are delta functions. An obvious way to enrich the basic MCM model so that it is more realistic is to allow the f_v to be conditional distributions.

5. Another way to enrich the basic MCM model is to have a special machine that serves as a **ledger**. This would allow us to model social institutions like the library for Ashurbanipal, or the clearing house of receipts in many ancient Levant societies — or the internet for modern society. Going further afield to consider different METs, if we model a single eukaryotic cell as an MCM (with the machines being organelles), the ledger could be the (epi)genome.

   In whatever context, a ledger machine's state space would be a tape (as in conventional TM theory), which may either be infinite, or finite (as in a genome) or finite and growing with the number of iterations, as in a library.

   In addition, what could make the ledger particularly useful in our model of a society agent is that we could choose a message graph such that (almost) all non-ledger machines would be able to send messages to and receive messages from the ledger machine. The effect on this ledger machine of the messages it receives from any other machine in a timestep is equivalent to that other machine writing to a specific location on the ledger machine's tape, and / or sending requests to read the contents of a specific location in the ledger machine. The state of the ledger machine would change appropriately in response to write messages. The messages sent by this ledger machine would then be different for each of the other machines, given by the ledger's response to the read request just made by those other machines (if any).

6. Similarly to a ledger, we could also have various machines that represent stores of material resources of various sorts. We could then have a message to or from the machine enforce conservation of the associated

good. For example, suppose a laborer machine v mines a single unit of some material resource from the environment at the beginning of an iteration. In the MCM framework, that means that at the beginning of the iteration, one component j of the observation vector, $e(j)$ of the society MCM has the value 1, specifying that single unit of that resource that was mined. In the first timestep, machine v increases some associated component i of their state vector, $x_v(i)$ by 1. For them to then send that unit of the resource to a storage machine v' would mean that in the associated timestep they send a message with the value 1 to v' and decrease $x_v(i)$ by 1 at the same time. Then in the next timestep, the value of $x_{v'}$ increases by 1. If at some subsequent timestep some other machine w needs to use that resource, a message can be sent from v' to w accompanied by a drop of 1 in the value of $x_{v'}$, and then the associated component of x_w increases by 1 when they receive that message in the next timestep.

More generally, production functions (in the standard economics sense) could be modeled as resource machines that send messages directly one another. The update function of a resource machine, setting the new value of its state based on the messages it just received and its previous state, would just be the associated production function we want to model with that machine.

7. Yet another way to enrich the model is to restrict the possible form of the message graph of an agent so that some nodes have far greater maximal fan-out than others. With that extension, introducing a new communication technology in a society agent would simply amount to introducing a new machine whose fan-out is far greater than that of the other machines in the agent.

Extending the model of interacting MCMs

The type of interaction between MCMs defined in Section 5 is an very pareddown, basic one. It can be enriched in many ways, to get a more complete (and therefore complicated) model.

1. One natural way to enrich the basic co-evolving MCMs model sketched above would be to have the noise in the two observation channels not be uniform and symmetric. For example, this would provide a way for the Society agent to pay especially close attention to some of the variables in the environment.

2. More generally, each machine v in an agent could get a different external input from the timestep 0 state of the the other agent, produced with a different observational distribution over values of ev conditioned on the joint state of the machines in the other agent. That set of machine-indexed observational channels would provide a **observation graph** for the agent.

   As an example, providing the Society agent with a observation graph would provide a way for some of the occupations in the Society agent to pay particularly close attention to the values of some subset of the variables in the Environment agent, e.g., due to geographical proximity. It may even be that some other occupations in the Society agent do not pay direct attention to *any* of the variables in the Environment.

3. Similarly, recall that the observation channels of the Environment agent are the same as the action channels of the Society agent. Accordingly, restricting which of the nodes in the Society agent are roots in the Environment agent's observation graph would allow some of the occupations in the Society agent to be more closely involved in acting directly on the Environment compared to the other occupations. This would be a way to capture the fact the members of different occupations observe (and affect) different variables in the environment (due to physical location of the members of the occupation, if nothing else).

4. Physical location can be quite important for how real human societies evolve from one timestep to the next, not just for how they can act on the environment at the end of an iteration. This raises the issue of how to represent the spatial locations of the machines in an agent. One natural way to do it would be by having the message graph of the Society MCM be a planar graph, with only nodes that (represent machines that we want to model as) physically close to one another directly connected to one another by that message graph.

5. Another important point to note is that in the basic model of the interaction of the Society and Environment agents in Section 5, there is a "harvest function" that specifies how the Society agent extracts free energy from the Environment, leading to new values of the parameters of the Society agent. However, there is no corresponding function that specifies how the Environment agent's state changes when free energy is extracted from it.

A natural extension of the basic model would capture this by providing special “resource storage” machine in the Environment agent, whose value specifies the amount of the resources whose harvesting provides free energy to the Society agent. There would then be a new step, just before the beginning of any iteration but after the end of the previous one. In that step the value of the resource storage machine in the Environment is reduced by the GFER generated by the harvest function of the Society agent (which in turn is based on the ending joint state of the two MCMs in the previous iteration).

6. Similarly, in a more elaborate model of the interaction of the Society and Environment agents, there might be resources which are transferred from the Environment agent to the Society agent but do not directly change the computation parameters or dynamics functions of the Society agent. These resources could be things like minerals, petrochemicals, building materials, timber, etc. To model a transfer of these resources, we would simply have machines in each agent that represent the total storage of the associated resource in that agent. At the beginning of an iteration, there could be a transfer of such resources between (the storage machines of those resources in) the two agents. Formally, this could be implemented without adding any more mathematical structure to the model of interacting MCMs above, simply by carefully incorporating “corresponding” dependencies in the observation graphs of the two MCMs, ensuring that the values in the associated machines in the MCMs changes under those two observation graphs in such a way that an increase in one results in a decrease in the other.

Discussion

The goal of this paper has been to identify the major gap that exists in the study of human social evolution. Namely, due to a series of contingencies in intellectual history, the study of social evolution has been disconnected from advances in computational complexity. It seems self-evident that there are advantages to bringing these two bodies of thought into conversation. It is much harder, of course, to bring this reunification about.

Here we have proposed a combined theoretical / empirical strategy to get at the issue. We think that our proxy, occupational specialization, has the potential to make concrete and visible the flow of information within a society and across evolutionary time. We have also proposed a formal, mathematical model of how a society computes and how its computational power changes over time. We have not, obviously, started the merger of our data and our model. That comes next, after the candid feedback of our august assembly, for which we express our advance gratitude.

There is at least one major body of work involving information theory to model the evolution of a species extracting resources from its physical environment. This body of work involves “Kelly gambling” and “bet-hedging”, and is discussed in Appendix C. It would be straightforward to incorporate that body of work into our framework, but we do not consider it here, for reasons of space. (See also [Guttenberg and Goldenfield, 2008] for a formal model that could maybe be usefully applied to investigate the METs.)

In contrast, we could not find any way that Ashby’s semi-formal “requisite variety” idea could enrich our framework. In fact, Baez and Aaronson (in separate blog posts) argue strongly that Ashby’s idea is either tautological or wrong, depending on how one carefully defines it.

Acknowledgment

We thank the Santa Fe Institute for support.

References

Adami, Christoph F. 2012. The use of information theory in evolutionary biology. *Annals of the New York Academy of Sciences* 1256: 49-65.

Aguilar, Wendy, Guillermo Santamaria-Bonfil, Tom Froese, and Carlos Gershenson. 2014. The past, present, and future of artificial life. *Frontiers in Robotics and AI* 1: 8.

Al-Hashimi, Hashim M. 2023. Turing, von Neumann, and the computational architecture of biological machines. *Proceedings of the National Academy of Sciences* 25 (120): e2220022120.

Arora, Sanjeev and Boaz Barak. 2009. *Computational complexity: a modern approach*. Cambridge University Press.

Arpe, Jan, Andreas Jakoby, and Maciej Liskiewicz. 2003. One-way communication complexity of symmetric boolean functions. Pages 158–170 in *Fundamentals of Computation Theory*, Springer.

Arroyo, Jose Ignacio, Beatriz Diez, Christopher P Kempes, Geoffrey B. West, and Pablo A Marquet. 2022. A general theory for temperature dependence in biology. *Proceedings of the National Academy of Sciences* 30 (119): e2119872119.

Babai, Laszlo, Anna Gal, Peter G Kimmel, and Satyanarayana V Lokam. 2003. Communication complexity of simultaneous messages. *SIAM Journal on Computing* 1 (33): 137–166.

Ball, Marshall and Tim Randolph. 2022. A note on the complexity of private simultaneous messages with many parties. 3rd Conference on Information-Theoretic Cryptography (ITC 2022), *Schloss Dagstuhl-Leibniz-Zentrum f ̈ur Informatik*.

Becker G. S. and K. M. Murphy. 1992. The Division of Labor, Coordination Costs, and Knowledge. *The Quarterly Journal of Economics* 4 (107): 1137–1160.

Bettencourt, Luis M. A., Horacio Samaniego, and Hyejin Youn. 2014. Professional diversity and the productivity of cities. *Scientific Reports* 4 (1): 5393.

Bettencourt, Luis M. A. 2013. The origins of scaling in cities. *Science* 340 (6139): 1438–1441.

Bettencourt, Luis M.A. 2020. Urban growth and the emergent statistics of cities. *Science advances* 6 (34): eaat8812.

Boyd, Robert. 2018. *A different kind of animal: how culture transformed our species*. University Center for Human Values series, University Press, Princeton. HOLLIS number: 990152078900203941.

Candogan, Ozan, Ishai Menache, Asuman Ozdaglar, and Pablo A Parrilo. 2011. Flows and decompositions of games: Harmonic and potential games. Mathematics of Operations Research 3 (36): 474–503.

Carneiro, Robert L. 1967. On the relationship between size of population and complexity of social organization, Southwestern Journal of Anthropology 3 (23): 234–243.

Carneiro, Robert L. 2018. *Evolutionism In Cultural Anthropology: A Critical History.* Routledge.

Chick, Garry. 1997. Cultural Complexity: The Concept and Its Measurement. *Cross-Cultural Research* 4 (31): 275–307.

Christian, David. 2004. *Maps of time: an introduction to big history, California world history library.* University of California Press, Berkeley.

Christian, David. 2017. Complexity, energy and information in Big History and human History. Pages 111–142 in Weller, R. Charles (ed.) *21st-Century narratives of world history*. Palgrave Macmillan.

Cover, T. and J. Thomas. 1991. *Elements of information theory*. Wiley-Interscience, New York.

Dargaj, Jakub and Jakob Grue Simonsen. 2020. A complete characterization of infinitely repeated two-player games having computable strategies with no computable best response under limit-of-means payoff. Pages 69-70 in *Proceedings of the 21st ACM Conference on Economics and Computation*.

Dargaj, Jakub and Jakob Grue Simonsen. 2022. Discounted repeated games having computable strategies with no computable best response under subgame-perfect equilibria. *ACM Transactions on Economics and Computation* 1 (10): 1–39.

Daskalakis, Constantinos, Paul W. Goldberg, and Christos H Papadimitriou. 2009. The complexity of computing a nash equilibrium. Communications of the ACM 2 (52): 89–97.

Dawkins, Richard. 1996. *The blind watchmaker: why the evidence of evolution reveals a universe without design*. Norton.

Denton, Trevor. 2004. Cultural Complexity Revisited. *Cross-Cultural Research* 1 (38): 3–26.

Donaldson-Matasci, Matina C. Carl T Bergstrom, and Michael Lachmann. 2010. The fitness value of information. *Oikos* 2 (119): 219–230.

Easterlin, Richard A. 1998. *Growth triumphant: the twenty-first century in historical perspective*, 1st pbk. ed. ed., Economics, cognition, and society. University of Michigan Press.

Enquist, Brian J James H Brown, and Geoffrey B West. 1998. Allometric scaling of plant energetics and population density. *Nature* 395 (6698): 163–165.

Enquist, Magnus Stefano Ghirlanda, and Kimmo Eriksson. 2011. Modelling the evolution and diversity of cumulative culture. *Philosophical Transactions of the Royal Society B: Biological Sciences* 366 (1563): 412–423.

Fischer, Orr Rotem Oshman and Uri Zwick. 2016. *Public vs. private randomness in simultaneous multi-party communication complexity, Structural Information and Communication Complexity*. Pages 60–74 in 23rd International Colloquium. Revised Selected Papers 23, Springer.

Forrester, Rochelle. 2009. *Guttman scale analysis and its use to explain cultural evolution and social change, How Change Happens: A Theory of Philosophy of History, Social Change and Cultural Evolution.* Best Publications Limited.

Galor, Oded. 2011. *Unified growth theory*. University Press, Princeton.

Gedeon, Peter. 2018. Social change or social evolution? Arguments for and against social progress in the sociological theory of evolution, Corvinus. *Journal of Sociology and Social Policy* 9 (1): 3–33.

Geertz, Clifford. 1973. *The interpretation of cultures: selected essays*. Harper torchbooks. TB 5043, Basic Books.

Goldenfeld, Nigel and Carl Woese. 2011. Life is Physics: Evolution as a Collective Phenomenon Far From Equilibrium. *Annual Review of Condensed Matter Physics* 2 (1): 375–399.

Guttenberg, Nicholas and Nigel Goldenfeld. 2008. Cascade of complexity in evolving predator-prey dynamics. *Physical review letters* 5 (100): 058102.

Hanson, J. W. , S. G. Ortman, and J. Lobo,. 2017. Urbanism and the division of labour in the Roman Empire. *Journal of The Royal Society Interface* 14 (136): 20170367.

Harper, Kyle. 2013. Culture, nature, and history: the case of ancient sexuality, *Comparative studies in society and history* 4 (55): 986–1016.

Hartle, Harrison, David Wolpert, Andrew Stier, Christopher P Kempes, and Gonzalo Manzano. 2024. *Work extraction with feedback control using limited resources*. arXiv preprint arXiv:2407.05507.

Håstad, Johan, Benjamin Rossman, Rocco A Servedio, and Li-Yang Tan. 2017. An average-case depth hierarchy theorem for boolean circuits. *Journal of the ACM* (JACM) 5 (64): 1–27.

Hausmann, Ricardo, C´esar A. Hidalgo, and Sebastian Bustos. 2013. *The atlas of economic complexity: mapping paths to prosperity.* The MIT Press.

Hayek, Friedrich August. 2013. The use of knowledge in society. Pages 27–38 in *Modern understandings of liberty and property*. Routledge.

Henrich, Joseph. 2004. Demography and cultural evolution: how adaptive cultural processes can produce maladaptive losses—the Tasmanian case. *American antiquity* 2 (69): 197–214.

Hilbert Martin and Priscila L´opez. 2011. The World's Technological Capacity to Store, Communicate, and Compute Information. *Science* 332 (6025): 60–65.

Hirvonen Juho and Jukka Suomela. 2021. Distributed algorithms 2020. Finland: Aalto University.

Jablonka Eva and Marion J Lamb. 2006. The evolution of information in the major Transitions. *Journal of theoretical biology* 2 (239): 236–246.

Jones, Charles I. 2002. Sources of us economic growth in a world of ideas. *American economic review* 1 (92): 220–239.

Jost, Jurgen. 2020. Biological information. *Theory in Biosciences* 4 (139): 361–370.

Jukna, Stasys et al. 2012. Boolean function complexity: advances and frontiers. Springer.

Kase, Vojtech, Petra Hermankova, and Adela Sobotkova. 2022. Division of labor, specialization and diversity in the ancient Roman cities: A quantitative approach to Latin epigraphy. *PLOS ONE* 6 (17): e0269869.

Kempes, Christopher P., Lawrence Wang, Jan P. Amend, John Doyle, and Tori Hoehler. 2016. Evolutionary tradeoffs in cellular composition across diverse bacteria. *The ISME Journal* 9 (10): 2145–2157.

Kolchinsky Artemy and David H. Wolpert. 2018. Semantic information, autonomous agency and non-equilibrium statistical physics. *Interface Focus* 6 (8): 20180041.

Koyama, Jared T., Mark Rubin. 2022. *How the world became rich: the historical origins of economic growth*. Polity Press.

Krakauer, David C. 2011. Darwinian demons, evolutionary complexity, and information maximization. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 3 (21): 037110.

Kuper, Adam. 1988. *The invention of primitive society: transformations of an illusion*. Routledge.

Kushilevitz, Eyal. 1997. Communication complexity. *Advances in Computers* 44: 331–360.

Langton, Christopher G. 1995. *Artificial life: An overview.* The MIT Press.

van Leeuwen, Marco H. D. 2002. *HISCO: Historical international standard classification of occupations.* Leuven University Press.

Li M. and Vitanyi P. 2008. *An introduction to kolmogorov complexity and its Applications*. Springer.

Lizier, Joseph T., Mikhail Prokopenko, and Albert Y Zomaya. 2014. A framework for the local information dynamics of distributed computation in complex systems,. Pages 115–158 in *Guided self-organization: inception*. Springer.

Marschak J. and R. Radner. 1972. *Economic theory of teams*. Yale University Press.

Mitchell, Melanie. 2009. *Complexity: A guided tour*. Oxford university press.

Mokyr, Joel. 1990. *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford University Press.

Monderer D. and L. S. Shapley. 1996. Potential games. *Games and Economic Behavior* 14: 124–143.

Moore Cristopher and Stephan Mertens. 2011. *The nature of computation*. Oxford University Press.

Morgan, Lewis Henry. 1877. *Ancient society; or, Researches in the lines of human progress from savagery, through barbarism to civilization*. H. Holt and Company.

Morris, Ian. 2010. *Why the West rules– for now: the patterns of history, and what they reveal about the future*. 1st ed. ed., Business book summary. Farrar, Straus and Giroux.

———. 2013. *The measure of civilization: how social development decides the fate of nations.* University Press.

———. 2022. Evolutionary history. *Evolutionary Psychology* 1(20): 14747049211068

Murdock George P. and Caterina Provost. 1973. *Measurement of cultural complexity*. Ethnology 4(12): 379–392.

Naroll, Raoul. 1956. A preliminary index of social development. *American anthropologist* 4 (58): 687–715.

Nisan N. and A. Ronen. 2001. Algorithmic mechanism design. *Games and Economic Behavior* 35: 166–196.

Douglass C. 1990. North, Institutions, institutional change and economic performance, Political economy of institutions and decisions, Cambridge University Press.

Papadimitriou, Christos. 2014. Algorithms, complexity, and the sciences. *Proceedings of the National Academy of Sciences* 45 (111): 15881–15887.

Peregrine, P. 2018. Toward a theory of recurrent social formations. *The Emergence of Premodern States: New Perspectives on the Development of Complex Societies* 2.

Peregrine, Peter N Carol R Ember, and Melvin Ember. 2004. Universal patterns in cultural evolution: An empirical analysis using guttman scaling. *American Anthropologist* 1 (106): 145–149.

Piccinini, Gualtiero. 2011. The physical church–turing thesis: Modest or bold? *The British Journal for the Philosophy of Science* 62(4):733-769.

Piccinini Gualtiero and Corey Maley. 2010. Computation in physical systems. *The Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/entries/computation-physicalsystems/

Pomeranz, Kenneth. 2000. *The great divergence: China, Europe, and the making of the modern world economy*. Princeton University Press.

Powell, Barry B. 2009. *Writing: theory and history of the technology of civilization*. Wiley-Blackwell, Chichester, U.K.

Rao Anup and Amir Yehudayoff. *Communication complexity: and applications*. Cambridge University Press, 2020.

Rivoire Olivier and Stanislas Leibler. 2011. The value of information for populations in varying environments. *Journal of Statistical Physics* 14: 1124–1166.

Romanowska, Iza, Colin D Wren, and Stefani A Crabtree. 2021. *Agent-based modeling for archaeology: simulating the complexity of societies*. SFI Press.

Romer, Paul M., 1990. Endogenous Technological Change. *Journal of Political Economy* 98 (5) Part 2: S71–S102.

Roughgarden, Tim. 2010. Algorithmic game theory. *Communications of the ACM* 53 (7): 78–86.

Sanderson, Stephen K. 1999. *Social transformations: a general theory of historical development.* Rowman & Littlefield Publishers, Lanham, Md.

Scheidel, Walter. 2019. *Escape from rome: the failure of empire and the road to prosperity.* Princeton University Press.

Schrodinger E. 1944. *What is life?* Cambridge University Press.

Schrodinger, Erwin. 2012. *What is life? With mind and matter and autobiographical sketches*. Cambridge university press.

Rogers Elman. 1971. *Service, Cultural evolutionism: theory in practice*. Holt, Rinehart and Winston.

Shin, Jaeweon Michael Holton Price, David H Wolpert, Hajime Shimao, Brendan Tracey, and Timothy A Kohler. 2020. Scale and information-processing thresholds in holocene social evolution. *Nature communications* 11 (1): 2394.

Sipper, Moshe. 1998. Fifty years of research on self-replication: An overview. *Artificial life* 4 (3): 237–257.

Sipser, Michael. 1996. Introduction to the theory of computation. *ACM Sigact News* 27 (1): 27–29.

Smail, Dan. 2005. In the grip of sacred history. *The American Historical Review* 110 (5): 1337–1361.

Smaldino Paul E. and Peter J. Richerson. 2013. Human Cumulative Cultural Evolution as a Form of Distributed Computation. Pages 979–992 in *Handbook of Human Computation* (Pietro Michelucci, ed.), Springer.

Smith, Adam. 1776. *An inquiry into the nature and causes of the wealth of nations*. Eighteenth century collections online, Printed for W. Strahan.

Smith, John Maynard. 1999. *The major transitions in evolution*. Perseus Books.

Smith, Sam, Damion Dooley, and Yongqun He. 2022. Foundational Development of an Occupation Ontology. *The Eighth Joint Ontology Workshops (JOWO'22),* Augus 15-19 2022, Jönköping University, Sweden.

Spier, Fred. 2015. *Big History and the Future of Humanity*. John Wiley & Sons.

Szathmary, Eors. 2015. Toward major evolutionary transitions theory 2.0, Proceedings of the National Academy of Sciences 112 (33): 10104–10111.

Szathmary Eors and John Maynard Smith. 1995. The major evolutionary transitionsю Nature 374 (6519): 227–232.

Touchette, Hugo. 2009. The large deviation approach to statistical mechanics. *Physics Reports* 478 (1-3): 1–69.

Trakhtenbrot, Boris A. 1984. A survey of russian approaches to perebor (bruteforce searches) algorithms. *Annals of the History of Computing* 6 (4): 384–400.

Tumer K. and D. H. Wolpert. 2000. Collective intelligence and Braess' paradox. *Proceedings of the Seventeenth National Conference on Artificial Intelligence (Austin, TX)*, 104–109.

Wolpert D. H. and K. Tumer. 2002. *Overview of collective intelligence, The Design and Analysis of Collectives.* Springer-Verlag.

Turchin, Peter, Thomas E. Currie, Harvey Whitehouse, Pieter Francois, Kevin Feeney, Daniel Mullins, Daniel Hoyer, Christina Collins, Stephanie Grohmann, Patrick Savage, Gavin Mendel-Gleason, Edward Turner, Agathe Dupeyron, Enrico Cioni, Jenny Reddish, Jill Levine, Greine Jordan, Eva Brandl, Alice Williams, Rudolf Cesaretti, Marta Krueger, Alessandro Ceccarelli, Joe Figliulo-Rosswurm, Po-Ju Tuan, Peter Peregrine, Arkadiusz Marciniak, Johannes Preiser-Kapeller, Nikolay Kradin, Andrey Korotayev, Alessio Palmisano, David Baker, Julye Bidmead, Peter Bol, David Christian, Connie Cook, Alan Covey, Gary Feinman, A´ rni Dan´ıel Ju´l´ıusson, Axel Kristinsson, John Miksic, Ruth Mostern, Cameron Petrie, Peter Rudiak-Gould, Barend Ter Haar, Vesna Wallace, Victor Mair, Liye Xie, John Baines, Elizabeth Bridges, Joseph Manning, Bruce Lockhart, Amy Bogaard, and Charles Spencer. 2018. Quantitative historical analysis uncovers a single dimension of complexity that structures global variation in human social organization. *Proceedings of the National Academy of Sciences* 115 (2): E144-E151

Tylor, Edward Burnett. 1881. *Anthropology: an introduction to the study of man and civilization*. Harvard College Library history of science project.

West, Geoffrey B James H Brown, and Brian J Enquist. 1997. A general model for the origin of allometric scaling laws in biology. *Science* 276 (5309): 122–126.

———. 2001. A general model for ontogenetic growth. *Nature* 413 (6856): 628–631.

White, Leslie A. 1959. *The evolution of culture; the development of civilization to the fall of Rome*. McGraw-Hill.

Wolfram, Stephen. 1984. Cellular automata as models of complexity. *Nature* 311 (5985): 419–424.

Wolpert D. H. and J. Lawson. 2002. Designing agent collectives for systems with markovian dynamics. *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems* (Bologna, Italy).

Wolpert, D. H., J. Sill, and K. Tumer. 2001. Reinforcement learning in distributed domains: Beyond team games. *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*: 819–824.

Wolpert, David, Jan Korbel, Christopher Lynn, Farita Tasnim, Joshua Grochow, Gulce Kardes¸, James Aimone, Vijay Balasubramanian, Eric De Giuli, David Doty, et al. 2024. Is stochastic thermodynamics the key to understanding the energy costs of computation? *PNAS* 121 (45): e2321112121

Wolpert David H. 2019. The stochastic thermodynamics of computation. *Journal of Physics A: Mathematical and Theoretical* 52 (19): 193001.

Xue BingKan and Stanislas Leibler. 2017. Bet hedging against demographic fluctuations. *Physical Review Letters* 119 (10): 108103.

Yang, Vicky Chuqiao Christopher P. Kempes, Hyejin Youn, Sidney Redner, and Geoffrey B.West. 2022. Scaling and the Universality of Function Diversity Across Human Organizations. https://arxiv.org/abs/2208.06487