INTELLIGENCE COMPLEXITY

Michael Swetnam
Robert Hummel
Charles Mueller
Paul Syers

POTOMAC INSTITUTE FOR POLICY STUDIES
901 N. Stuart St. Suite 1200
Arlington, VA 22203

901 N. Stuart St, Suite 1200
Arlington, VA, 22203
www.potomacinstitute.org
Telephone: 703.525.0770; Fax: 703.525.0299
Email: webmaster@potomacinstitute.org

The Potomac Institute for Policy Studies is an independent, 501(c)(3), not-for-profit public policy research institute. The Institute identifies and aggressively shepherds discussion on key science, technology, and national security issues facing our society. From these discussions and forums, we develop meaningful policy options and ensure their implementation at the intersection of business and government.

Cover image credits:
Thinker image: Alex Taliesen.
Fractal realms series images: www.shutterstock.com.

Potomac Institute Press

TABLE OF CONTENTS

ACKNOWLEDGEMENT

The theory of DIKW as a hierarchy of intelligence complexity and the concept that I = E x C are concepts of Michael Swetnam that were vastly improved and corrected by the assembled quartet. DIKW as a measurement of intelligence was formalized by Robert Hummel and has been submitted for a US Patent. The writing and final product are the works of Charles Mueller and Paul Syers.

PART 1

DIKW THEORY OF INTELLIGENCE COMPLEXITY

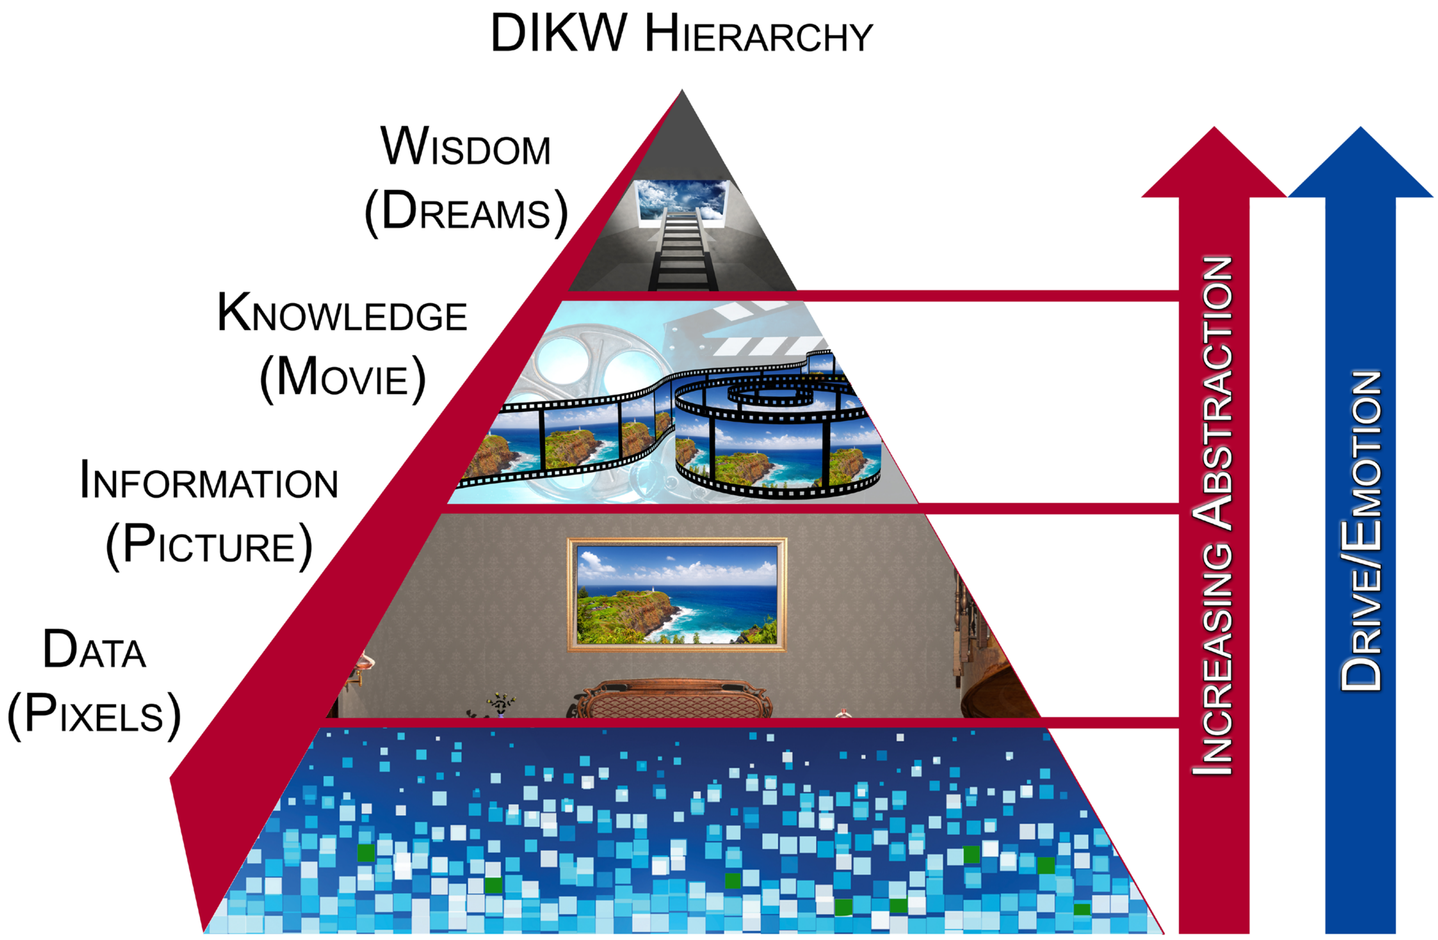

DIKW Explained

Today the world is connected like never before. Our actions, our locations and even our thoughts are aggregated and stored into large data sets that represent all we do in the digital environment. Companies that attempt to find patterns in these data sets are said to be working in the "Big Data" market. The notion of "Big Data" is about finding patterns in large data sets that can be abstracted into models that can lead to a more complete understanding of why such patterns exist. Companies care about this because it can be used to predict risk, human behavior, and even things like health. While humans are able to create these abstractions rather intuitively through reflection on their experiences in life, the companies interested in doing this use computers to find these trends based on our digital activity. Being able to identify patterns and build models to explain those patterns is, in a nutshell, how we view the idea of intelligence. This is where the notion of DIKW (Data, Information, Knowledge, Wisdom) finds its home because it can help provide a framework to understand intelligence, regardless of source, in terms of distinct levels of complex thinking (i.e., intelligence complexity).

The notion of DIKW more or less began when Russell Ackoff, a systems theorist, famously drew a distinction between data, information, knowledge, and wisdom, as a hierarchy of levels of abstraction of data. When Ackoff described the DIKW hierarchy in 1979, he used very general language to define the various components of the hierarchy. We hope that our descriptions sharpen these notions, although we admit that they may diverge from the original intention. In any case, for our purposes of attempting to describe the complexity of intelligence, something that could be used to motivate the development of higher order machine intelligence, as well as devising tests to know when intelligence at higher complexity levels has been achieved, we will need very precise definitions that include mathematical formulations. Accordingly, what follows is our attempt at a technical formulation based on mathematical frameworks.

These are not formulas, but rather frameworks into which particular instances of DIKW could fit.

For simplicity, we discuss intelligence as a generic system, almost in a thermodynamic sense. The generic intelligent system is capable of measuring signals from its surroundings, storing the signals as data, processing them, and generating some output in response to them. These output levels involve larger and larger amounts of abstraction from the direct sensory inputs from a system. The increasing levels correspond to an ability to express representations of the system using smaller and more complex methods, indicating higher levels of understanding. Instead of arbitrarily naming the levels, we follow Ackoff's original formulation and use the terms "information," "knowledge," and "wisdom," and the term "data" as the foundation. Thus we begin with data.[1]

Data

Data level intelligence complexity provides the most basic and least dynamic understanding about a system, by simply recording and storing, and recalling sensory inputs from the system.

In the general sense, data are just the signals coming into a system that can be detected, stored and processed. Abstractly, data is a process that results in a set of numbers or values that are the measurements or recordings of sensory input. One can think of it as a collection of bits, numbers, or recorded "things" that have associations to their source, which is the surrounding system. It is this evidence, whether observed or measured, that forms the building blocks of the DIKW structure. Again, by a system, we mean anything that is reasonably independent of other things in terms of functionality and interface. It can be a simple closed system that has a few inputs and outputs, or it can refer to something as complex as a solar system.

1. Russell L. Ackoff, "From Data to Wisdom," *Journal of Applied Systems Analysis* 16 (1989): 3-9.

Thus, data is a process that records details of a system and outputs measurements. Mathematicians talk about a "representation" of the system, although the representation may describe only a tiny part of the entire system, and the measured values may be approximate values. The realization of the measurements can include identifiers such as a time stamp.

Data is not necessarily "truth." For example, temperature readings of a system might provide measurements that are noisy and thus estimates of a "true temperature"; they constitute the fact that the data given by the sensor system (the thermometer) is recorded at a particular point in time and subject to particular conditions. When digitized, the data can be stored using bits, together with information about the time and source, and potentially additional information about error brackets, bounds, accuracy, and other parameters.

Data can also be text, or descriptions, in an unstructured format. Often, the collection of data is not particularly well organized. However, one of the main uses of computers over the past few decades along with the transition from analog to digital data has been an ability to automatically organize data in ways that are more useful for subsequent analysis.

Based on this definition of data, it should more clearly follow why we interpret the complexity of data level intelligence as being limited to the direct, basic operations of recording, storing, and accessing individual pieces of data.

Information

Information results from the application of a process to a dataset (or multiple datasets), which establishes a relationship among various pieces data, such as a correlation or average. Information is therefore derived from data. It is not based on direct measurements, but rather, produces new understanding based on relationships among data elements. It adds new meaning.

In a mathematical sense, such relationships might be provided by a regression of numerical values, or a retrieval of a record based on specific criteria, or a statistical database operation that combines more than one datum. For text data, information might be a summary or synopsis developed from the data, or an explanation that comes from combining text data with other data or other information. Often information comes from finding relationships between different combined data sets.

Information is created, since a body of data has to be absorbed. Information is about patterns/trends in a dataset, and not about a single piece of data. Thus, information describes constraints on the data, at least in an approximate sense. As a result, information includes a certain level of predictive power, gained from understanding data. Whereas data has no interpretative power, and hence cannot be said to contain any intelligence, information begins the process of moving up an intelligence hierarchy, by virtue of examining a body of data in the context of a question or other data.

Mathematically, information involves functions applied to a dataset. The information is the collection of functions together with the set of data points, or output values. For example, the average value of the data set is the result of a function that evaluates the average. Both pieces (the averaging function and the data set being averaged) provide less understanding of the system than the information of the average of the data set. Other functions might perform a linear regression and provide the parameters of that regression; another function might describe the data as following an approximate exponential growth pattern. It is the functional that codifies the information, which describes the trends in the data, or a specific operation applied to the data together with the resultant value.

Information retrieval, for example, happens based on looking at the entire database, together with the resultant extracted results. Most importantly, however, information is at a higher level of complexity than data, and can be distinguished by the fact that it specifies constraints, patterns, or statistics about data.

Knowledge

Knowledge increases the predictive power of information in an essential way. Knowledge involves the formulation of a model, which extrapolates beyond the experiences in the observed data, by providing a causal explanation of the data. The data provides the observables (i.e., measurements) taken from the system, but the model of the system attempts to explain how the system works, and thus should be consistent with the data, but also extrapolate from it. Accordingly, it predicts what the data might look like in other kinds of situations.

Knowledge is distinct from information in that the model of understanding of the system can hypothesize causation and underlying structure to explain the behavior. The model is more complex than a simple functional relationship. It relates a larger number of variables. A linear regression of data, while a primitive model that includes a few parameters, does not explain causation at any level, since the relationship between the data elements is correlative rather than causative. There are many examples of correlation that have nothing to do with causation.[2] Correlation provides a global structure, but not an underlying constituent structure. Models provide an understanding of underlying structures and the ability to predict data that have never been experienced.

As discussed previously, the world of "Big Data" is focused on developing machines (i.e. computers) that can obtain more complex levels of intelligence. The key to achieving this is a machine's ability to create models. Accordingly, we should consider the constituent components of a model. When we speak of models, we generally mean something that explains how inputs, or the state of the system, are related to the predicted outputs or progression of the system. Further, a model ideally predicts the behavior of the system in cases that extrapolate from observed data, or observed

---

2. Lee Falin, "The Truth (and Lies) of Correlation vs. Causation," *Scientific American*, last modified October 2, 2013, http://www.scientificamerican.com/article/correlation-vs-causation.

experience. In other words, we speak of models that go beyond being a set of correlations to identify causation. A model is often the hypothesis in the use of the scientific method, and is validated through experimentation that verifies the predictions outside of the range of existing experience.

Models must be useful for predictions, particularly beyond observed phenomena. But models are often refined as more data becomes available and experiments show discrepancies, however minor, from the existing model. While we tend to want to think of a successful model as being "truth," it is in fact an approximation, up until the time that it is refined so as to provide a better approximation. For our purposes, however, a successful model is useful. It provides predictions that can be used to understand why things behave the way they do, and to predict how things might behave in other circumstances. To operate at the knowledge level, a model needs to be useful and it needs to be able to iteratively change when it accumulates new data.

Knowledge in the end is an abstraction of the patterns and trends contained within information. It is the bigger idea contained within several different sources of information. This bigger idea helps explain why the information exists as it does, it provides causation, a model to help explain how the data is organized into the patterns and trends observed.

Wisdom

Beyond knowledge there is wisdom on the hierarchy. Wisdom is an abstraction of knowledge, which is itself an abstraction of information, which is itself an abstraction of data. In this sense, it is like a third degree abstraction of data. We posit that wisdom necessarily involves a functional model whose elements are bodies of knowledge, which is to say a model of models. The first-order models are the elements of knowledge; the model that assembles those models is a second-order, or meta-model, with far greater predictive power. As with knowledge, a meta-model should be able to extrapolate beyond the previously observed experiences.

However, meta-models allow for "what-if" experiments, and thus wisdom goes beyond extrapolation. Wisdom comes from sufficiently broad bodies of knowledge such that we might be able to postulate changes. Those changes might suggest we could influence or modify the generation of data or information through manipulation of a manageable set of input parameters and through wisdom, understand the likely impact.

Wisdom involves one or more meta-models, and invokes multiple knowledge-based models in order to provide sophisticated simulations and explanations of behavior. It is distinguished from knowledge by the use of multiple models, and further extrapolation to events that are not included in the information base. That is, the meta-model may involve positing a sequence of events and predicting the resulting data outputs. Importantly, wisdom involves a notion that the observer can control the outcome by manipulating events. Wisdom is the most dynamic of the levels of intelligence complexity. It cannot only adapt multiple models, but it can create entirely new ones for testing and incorporation into the meta-model.

In wisdom, the causation models are bound to the meta-model at prediction time, which is to say that if one of the models changes, then the result of the meta-model changes. That may seem obvious, but the important distinction is the difference between a meta-model and a model. If the meta-model is compiled to become a large model, then it is simply another model and can be considered additional knowledge. That new model is independent of the constituent models: They have been bound to the meta-model at the compilation time, and not at the "run time," to use a computer science expression. The meta-model, which is wisdom, uses knowledge models in such a way that if the knowledge models are dynamically updated or envisioned as different models, then the meta-model automatically uses the updated model. In this way, the meta-model can consider what we might term as "alternate realities."

For a machine, in order for wisdom to be used to influence outcomes, there is a separation of the input variables of the final program into variables that are observed and variables that can be controlled. Further, the variables might have a time sequencing requirement, or particular time differentials that are specified and intended. In wisdom, we use the prediction capabilities to seek goals by manipulating the controllable variables in order to obtain desirable predicted results.

As described above, wisdom is built upon multiple models of knowledge. However, because wisdom uses knowledge models at the time wisdom is invoked, those knowledge models are fungible. In this way, wisdom can play "what if" games. Wisdom can instantiate knowledge models that are not real, or not based on gathered information. Instead, wisdom can rely on hypothesized knowledge models; to "dream" (if you will) about alternative possibilities, even if the knowledge is not real. Wisdom allows us to develop theories about how things might work in an alternate universe, a different planet, or a different time in history.

Thus, wisdom allows for creativity in the application of the meta-model. Creativity is not present in information. In knowledge, we have more creativity in hypothesizing causality, but that creativity is limited in scope and by the need to validate the causal models. Wisdom is inherently creative, it is imaginative since the knowledge that is invoked as inputs need not be reality.

In wisdom, the first-order models are the elements of knowledge; the meta-model that assembles those models is a second-order, with greater predictive power and an ability to speculate on alternative first-order models. We can potentially consider meta-meta models, or third order (or higher) models that are built on top of lower-order models. In this way, wisdom itself can have multiple discrete levels of intelligence defined by their complexity. For this exposition, we will lump all such levels into a single category of "wisdom."

DIKW TEST

How will we know when another system has achieved intelligence at the knowledge or wisdom complexity levels when we have it? As we have asserted, the process of achieving that level of intelligence complexity should be transferable to multiple application domains.

For specific applications we can match the alleged system with intelligence against the capabilities that signify information, knowledge, or wisdom level intelligence complexity. However, there may be systems with other application domains, or a mixture of application domains that will have capabilities matching none of the patterns we have suggested for systems with intelligence within a specific domain. Accordingly, we consider how to determine the level of intelligence complexity of a proposed system against a DIKW hierarchy.

Within each level, the level of performance can be assessed through its predictive power. We have already noted that knowledge is subject to the scientific method in validating hypotheses. That is, a model that is generated by knowledge can be assessed through tests that validate its predictive capabilities.

Measuring the performance of wisdom is complicated. It is about measuring creativity. Wisdom comes from succeeding with predictions. Ultimately it will be the sophistication of the meta-model, and the utility of that model for influencing future data that provides the true measure of wisdom.

Systems and Methods for Determining Machine Intelligence[3]

> We propose a 20-questions format for determining the intelligence complexity level of a system, whether it is at a D, I, K, or W Level. This test does not provide for a measure within the level, but just the highest level of attainment for a system capable of intelligence. Further, we

3. Systems and Methods for Determining Machine Intelligence is being patented and has a provisional patent number 15/198,942.

can insist that the system pass certain questions at a given level, and that for other questions, one or more of a group must be satisfied. By assigning points to each question, and then scoring the system at each given level, we can decide if the system has truly attained that level of intelligence complexity, while still allowing for some ambiguity and "partial credit" in some of the questions.

An intelligent system can operate at a data level, an information level, a knowledge level, or a wisdom level (or none of the above), which we call D, I, K, and W Levels. If it operates at an I Level, then it also either operates at or uses a system that operates at the D Level. Similarly, a K Level system is also, or uses, an I and D Level system, and a W Level system is also, or uses, K, I, and D Level systems.

For a system X that purports to operate at the D Level

1. Does X receive inputs that are measurements (data)? (20 points)
2. Does X insert those measurements into a store of data? (20 points)
3. Does that store of data have permanence, such that it can be appended or reviewed later? (20 points)
4. Does X permit subsequent use of that store of data? (20 points)
5. Can database operations be executed on that store of data? (20 points)

If a system scores 90 or higher (allowing for some ambiguity in the scores for answers to questions), then X is at least a D Level system.

For a system X that purports to operate at the I Level

1. Does X permit queries that request information? (25 points)
2. Does X access multiple elements of a data store in order to answer the query? (25 points)
3. Does X find trends in the data and output information about those trends? (10 points)
4. Does X find statistics concerning the data and use those statistics to provide information? (10 points)

5. Does X correlate data across the data store, or find correlations among the elements in the data store? (10 points)
6. Can X predict data that would be measured for a system that interpolates between states of the system for which data has been collected? (10 points)
7. Does X combine data from more than one database? (10 points)

If X scores greater than 80, then it is at least an I Level system.

For a system X that purports to operate at the K Level

1. Does X ingest or build information about a system? (20 points)
2. Does X build a model of that system, such that the model depends on the information that X receives? (20 pints)
3. Does the model include a model of causality that explains how the system works or evolves in response to its inputs? (10 points)
4. Can the model provide useful predictions of information about the system that it models? (10 points)
5. Does the model include a set of values that corresponds to a notion of the state of the system that is being modeled? (10 points)
6. Does the model explain most of the information that is provided about the system? (10 points)
7. Does the model permit the prediction of information that extrapolates from the observed behavior of the system on which the input information was based? (10 points)
8. Does the model provide information about the structure of the system, including elements that cannot be directly observed and are thus not part of the input information? (5 points)
9. Can X build models about different systems, based on input information about each such system? (5 points)

If X scores greater than 80, then it is at least a K Level system.

For a system X that purports to operate at the W Level

1. Does X ingest multiple models that represent a compound system, where each one models either all or part of the system (i.e., a subsystem)? (20 points)
2. Does X build a model (a meta-model) of a system that varies if any of the ingested models vary? (20 points)
3. Can X change one or more of the ingested models, to thereby change the output meta-model (in a what-if experiment)? (10 points)
4. Does X use the meta-model to explore possible states of the modeled system, under various hypothetical circumstances (states)? (10 points)
5. Does X use the meta-model to explore possible states of the modeled system by varying ingested models? (10 points)
6. Does X use the meta-model to explore possible states and attempt to maximize a metric applied to the information provided by the meta-model? (10 points)
7. Does X provide information about how the system might be changed so as to provide different (and better) states, according to some metric? (10 points)
8. If so, is that information actionable, in that controllable parameters of the system could be changed so as to conform to the different and better state of the system, as predicted by the meta-model? (10 points)

If X scores greater than 80, then it is a W Level system.

We have proposed that there are discrete complexity levels of intelligence. However, within a given level, it is possible (indeed, probable) that there are measurable degrees (or a continuum of degrees) of intelligence complexity at that level. It is not likely that tests exist to measure the degree of intelligence complexity at the information, knowledge, or wisdom level, although it is possible that there are surrogate approaches.

The Utility of A DIKW Test

The test for intelligence complexity, such as the scoring test in the previous section, yields a discrete measure of the level of complexity at which an intelligent system can process. This will, for example, enable developers of machine intelligence to assess the degree of success at climbing the DIKW hierarchy, and afford researchers of other forms of intelligence a similar tool.

Specifically for those focused on creating a machine with higher levels of intelligence complexity, this tool provides a different measure of success – not one focused on the monetary value of the capability of an algorithm or machine in the marketplace. For example, Google is able to sell advertisements based on its ability to collect information and use that information to the benefit of their clients. We have seen that monetary gain and many other measures of machine utility tend to argue for better information level machines.

But from the standpoint of research, we aspire toward higher goals that better explain nature. Companies that develop and market services that provide machine intelligence will have greater marketability if they can provably claim that their machine has knowledge level intelligence complexity, whereas competitors are at information level intelligence complexity. Researchers who attempt to understand if other species on Earth (or beyond) are capable of human-like intelligence can use this framework as a guide in the quest for a true universal theory of intelligence. By defining the difference between levels of intelligence complexity for any system, we set aim points for the development and discovery of successively higher complexity levels of intelligent systems.

The test described in the previous section seeks performance that can reason about causality and provide guidance to influence outcomes, with greater predictive fidelity. In addition to providing greater functionality, this aim point will help us understand the nature of intelligence, as opposed to using tricks to mimic intelligence.

On a final point, in the realm of machine intelligence, an ability to provably attain wisdom level intelligence will offer far greater capabilities. Beyond the bragging rights, knowledge and wisdom level abilities would be able to achieve better predictions, greater marketability, and indeed give the user the ability to generate new knowledge, or to use higher levels of wisdom.

Furthermore, it is likely that this test can be automated. Modern compiler technology can be used to examine the code of a machine, and scripts can be used to execute sample runs. At issue would be whether the machine accesses the database in ways that mix data from multiple sources, whether the machine creates predictive models with a state space that extrapolates from the information in the database, and whether the machine is able to hypothesize novel knowledge bases and thus consider ways to influence outcomes. While the test provides opportunities for studying and understanding the complexity of intelligence writ large, its most direct and immediate application clearly seems to be directed towards the drive for machines with higher levels of intelligence.

An Example of DIKW in Action

With the different complexity levels of data, information, knowledge and wisdom better defined, let us look at two examples to illustrate the different types of intelligence complexity in operation. First we will look at how a human DIKW system operates and then compare this to current capabilities of a machine DIKW system (i.e., computer).

From D-W: Human Intelligence

Humans are the only observable intelligent system known to be capable to wisdom level thinking. As such, any descriptions we have about an intelligent system that can perform at all levels of the DIKW hierarchy are limited to a discussion of humans.

Humans are born with the potential to develop wisdom, but they are not born with wisdom. Humans develop wisdom only over time and there is nothing guaranteed they will ever actually obtain it. The path to obtain wisdom begins before a child is even born. As the body and mind are developing, data is being delivered into a young infant's brain through all of its newly developed measuring tools, its senses. At this point, a child is only capable of data level thinking. It has no ability to yet make sense of all the different signals coming into its brain.

During the first years of a child's life, the brain begins to notice trends and patterns within its experiences. It learns to speak by recognizing how the sounds it hears match with the facial expressions it sees and the sensations it feels with its own ability to create such sounds. It forms these patterns by storing its experiences in "memory" and "recalling" them at different times to compare them to existing experiences. A child at this point becomes capable of intelligence at a complexity level of information. It can mimic and engage in meaningful ways with its environment, but the child doesn't really understand why anything is happening around it or within it. It has no knowledge.

Farther down the line we begin to see the emergence of a human intelligent systems pursuit of knowledge. It begins with questions of "Why?" Once a child begins to question the patterns and trends it recognizes and remembers, it begins to go through the iterative process of building models to explain all the data and information it observes. The biggest problem human intelligent systems have once they begin asking why is not having at least some answer to explain things. The answer doesn't have to be right, it just has to be useful and at this stage in the game, any answer is better than no answer. As such, the child's mind goes from thinking the Tooth Fairy, Easter Bunny and Santa Clause sneak into their house in the middle of the night to deliver gifts to realizing the presents are all part of a human tradition based on letting the ones we love know we care. Through schooling and their own experiences, driven by emotion that fuels their curiosities and imagination, humans develop a vast array of models to explain the world around them.

With discipline and a commitment to practicing a structured approach to the pursuit of knowledge, a human system can begin to connect the models it has to explain all the things around it. It begins to understand the complexity of human nature, the laws of nature, and how it is all connected. They develop a real model of models and that is what we call wisdom. The domain of wisdom might only be applicable to certain experiences, like love, or certain disciplines, like physics, but there can be even greater levels of wisdom that begin to understand how things like love and physics are related.

It is through the iterative process of gathering data (experiences), identifying the patterns within it (information), creating models to explain it (knowledge), and creating even greater models to validate and explain the existence of these models (wisdom) that the human intelligence system operates; it is a process of creating higher level abstractions. Many human systems never achieve wisdom, and many of them develop knowledge level models that do not factually explain the world around them, but are useful to their way of life nonetheless. The human intelligence system reminds us that just because one has the potential to develop wisdom, it does not mean they actually have and practice wisdom. The DIKW test described in this report measures this potential, not its application.

From D-W: Humans & Climate

The fact that humans can understand and model the climate is an example of their potential to achieve wisdom level intelligence. The drive to understand the climate was likely one of the first real pressures on humans to evolve a greater intelligence. The ability to understand and predict how the weather might change could literally at times mean the difference between life and death.

It all starts with collecting data. Humans can collect data using instruments like thermometers, barometers, anemometers, hygrometer, and rain gauges. By collecting a lot of data over time about how a particular region on this planet is changing with respect to things like temperature, humidity, and rainfall, humans can process this data into information. A human intelligent system can collect a region's data throughout the course of a year and learn things like average temperature and rainfall, patterns in pressure changes, etc. for that region. Collecting it long enough even allows for the understanding of long-term patterns in how these data measurements change.

With this information, human intelligence can anticipate things like the ranges of temperatures to expect for a given time of the year. They can plan for certain times of the year where there is more rain and other times when it is more dry. They can prepare for the times of the year where the storms are the strongest or the winters the coldest. This ability to create a predicative capability from this information is the jump from information to knowledge. It requires an ability to create a model about the local region being studied that provides some sort of causative mechanism for why the information patterns/trends they observe exist. We call this model weather. The weather associated with the seasons are abstractions of the unique patterns we relate to particular times of the year. When data starts to show X, leading to the pattern Y, we understand it to be season Z.

Although, if we leave it here, we still lack a lot of predictive power. Our models cannot explain the anomalies, the hot days in the winter or the cold days in the summer. It can't help us understand why it rains at all. It can't help us understand how to look at any one day and be certain of what tomorrow will bring. For that we need a better model. If we expand from the region with local knowledge into a larger model that connects other regional models all over the planet, we begin to see and create a model that describes how the weather at any place in the world exists. In this sense, we create a climate model, which is approaching the wisdom level of intelligence complexity.

True wisdom at the level of climate would allow one to predict weather that hasn't happened yet, to predict how changing some variable that is part of the climate model will change other variables and in that sense, create new data; it allows "what-if" thinking where one can imagine how to create new data, information, or knowledge. With a wisdom level model, one can imagine how they want the climate to change and perform the necessary adjustments to create the climate, the world they see in their mind.

Currently humans are on the verge of true wisdom with respect to climate, but true wisdom is hard to obtain. However, even if our wisdom level models of climate are not perfect, they are useful, helping demonstrate that humans are the only intelligence systems known that can achieve this level of intelligence.

From D-I: Machine Intelligence

Netflix has employed Big Data Analytics to help influence its business decisions, and a notable example of its success in using this method, the creation of the TV show *House of Cards,* involves results of intelligent outputs at two of the four levels of the DIKW hierarchy. The point of this example is to demonstrate the current capabilities and limitations of machine intelligence.

One of Netflix's activities, in addition to providing movies and television shows to stream online, is to collect massive amounts of data about its users: not just what movies and shows they like, but also what day of the week and what time of day they watch programming, whether they watch an entire program, what kind of device they use to watch it, demographic details, and even geographic location details. This step in the process is all at the data level. It is logging bits of data (e.g., this viewer liked *The Goonies,* didn't like *Transformers,* watched first 15 minutes of episode 1 of *Breaking Bad* on a mobile device in Poughkeepsie, NY, etc.) At the data level, someone or some machine at Netflix can retrieve those data points mapped to a particular user.

The machines used by Netflix to perform the Big Data Analytics on this data go the next step of recognizing correlations, which is information level thinking. These Big Data algorithms demonstrated that there was a correlation between Kevin Spacey being in a film/show and the average rating of that show being high, as well as the same correlation with the British show *House of Cards*. Another set of correlations was identified regarding the director David Fincher. One was that a high percentage of users that watched one of David Fincher's works also watched all of his works and a second correlation was found between David Fincher's works and a high average rating. There was also the correlation discovered that a high percentage of people who rated Kevin Spacey's works highly also rated David Fincher's works highly, and rated *House of Cards* highly (this correlation amongst more than two variables is known as a strong correlation). Netflix's Big Data algorithms applied this functional of co-occurrence to the data sets Netflix had gathered and recognized the common details of the groupings that displayed this co-occurrence. It was correlative thinking and therefore information level intelligence complexity.

For machine intelligence though, this is the end of the line. The success of the show *House of Cards* from here can only be attributed to human intelligence. Yes, the Big Data Analytics made things easier for humans, but that is nothing new for technology. The knowledge and wisdom level thinking required to determine that a show based on a successful British show, directed by David Fincher and starring Kevin Spacey would generate more revenue for the company could not be determined by Big Data Analytics; it requires more complex thinking. These current systems lack the necessary motivations to pursue the kind of creative thoughts necessary to achieve the things that human intelligence can. As such, we examine this concept of a necessary system to drive an intelligent system up the DIKW hierarchy in the following section.

PART 2

THEORY OF EMOTION

The Driver of Intelligence

What drives and motivates an intelligent system to collect data, trend it into information, build the models of knowledge, etc.? Without a motive force to drive such a system it would not execute the processes that define DIKW.

If intelligence is a universal concept, than a relationship must exist between intelligence and the force that drives its progression. Given that the human intelligence system is the only observable intelligent system we know capable of traversing the DIKW spectrum, it is therefore the only system for which we can examine the force that drives the progression up the DIKW hierarchy.

Humans are driven by emotion. We are motivated to observe our universe (collect data). We try to build models of understanding more because it feels good to solve problems, to know how to do things more efficiently and to have answers to difficult questions. We are emotional creatures and thus it seems rather intuitive that there must exist some sort of relationship between this concept we call emotion and intelligence.

We describe in this section an idea that intelligence is driven by emotion, meaning an intelligent system requires an equivalent to human emotion in order achieve higher levels of intelligence. In humans, emotion is the source of our curiosity, creativity and other motivating forces to try new things, have new thoughts and experience more. These drivers provide the motivation for an intelligent system to evolve – to become more intelligent. Without a driver, the system lacks the motivation and energy to climb the DIKIW hierarchy.

I = E x C

The theory presented in this report regarding DIKW is not one that attempts to explain the very nature of intelligence. Rather it is a theory that explains how one measures the complexity of a system in terms of intelligence (i.e., intelligence complexity). It provides a measure of a system's ability to reach defined levels of intelligence. As such, any system that can demonstrate and be measured at one of the levels defined by DIKW can be said to be an intelligent system capable of intelligence processing at that level. Up to this point, this discussion has not addressed what an intelligent system requires to actually process and improve its ability to process (i.e., climb DIKW hierarchy) at a particular level. We argue that to actually process autonomously at a level of DIKW requires a source of motivation, a drive to actually do so.

We propose this drive is a force commonly referred to as Emotion (E). Further, we put forward the idea their exists a universal relationship between Intelligence (I),[4] Emotion (E), and the Intelligence Complexity (C) level (i.e., DIKW level) of an intelligent system that can be stated as:

$$I = E \times C$$

This general framework simply states that in order for a system to demonstrate intelligence (I), it must have a motivational driving force (E) and be capable of some level of complex processing (C) described by DIKW. It does not explain yet, why and how such a relationship might exist.

What follows is an attempt to further explain this relationship. We do this by examining the only known intelligent system capable of demonstrating the highest levels of intelligence complexity (i.e., DIKW) and is driven by some emotional force (i.e. the human system). This justification centers on a thermodynamic argument that attempts to explain the human intelligent system in terms of complexity, efficiency and entropy.

4. Here, Intelligence (I) is not a direct measure of how high up on the DIKW hierarchy an intelligent system is, but rather is an abstraction of the energetic output an intelligent system creates; it is the response of an intelligent system to some stimuli that demonstrates the system has intelligence.

I = E x C: The Human System

A mature human often has an intelligence system that is complex, ordered, and capable of traversing the DIKW hierarchy.

Emotion is then the force or the fuel that is used to build an intelligent system and the responses it creates to the stimuli it can detect, store and process. The more stimuli and ways a particular stimulus can be detected and interpreted, the more "energy," in the form of intelligent work, that particular stimuli can provide the system. Thus one would predict that systems capable of higher levels of intelligent thinking would also require more emotion to drive this process.

In order to increase the efficiency of turning emotion (E) into intelligence (I), the human system evolves to higher levels of complexity.[5] We describe part of the way the human system achieves this through DIKW, as it is a measure of the systems complexity in terms of intelligence (C). The higher up an intelligent system finds itself on DIKW, the more complex "thoughts" that system can have. These complex thoughts are what we understand as curiosity and creativity. They are the ability to question and create abstractions that help simplify the world the intelligent system is a part of. As such, the more complex thinking an intelligent system is capable of, the more "efficient" that system ends up being at converting a particular amount of "emotion" into an intelligent response.

Based on this framework, it follows that for two intelligent systems that have the same amount of emotion (E), the one capable of more complex thinking (i.e. higher DIKW; C) will produce the more intelligent response. This would be seen as a more efficient utilization of the energy it has available and its actual manifestation could be a number of things. It could be that the intelligent system uses that available emotion to generate an

5. Shahar Dolev and Avshalom Elitzur, "Biology and Thermodynamics: Seemingly-Opposite Phenomena in Search of A Unified Paradigm," *The Einstein Quarterly: Journal of Biology and Medicine* 15 (1998).

actual response (i.e., some answer to a query), create a better model of its existence/environment,[6] and/or use it to do "housekeeping" (i.e., repair and maintain the system). This would be opposed to the other intelligent system that can only convert the energy into some partiality of this work. The point being, the greater the intelligence complexity term, the more efficient the intelligent system is at converting the emotional energy into an intelligent response.

In summary, our theory regarding the relationship between emotion and intelligence hinges on the principal that intelligent systems are inherently complex and require a motive force to build and maintain its complex, ordered nature. The motive force we are proposing is something we humans understand as emotion and we argue this is the force used to drive an intelligent system to "learn" and "grow." The ability of an intelligent system to convert emotional energy into intelligent work is controlled by an intelligence complexity factor, C. This intelligence complexity factor is a measure of the system's ability to have complex thoughts and is directly related to the intelligence complexity described by the DIKW hierarchy, meaning there should be a complexity factor term characteristic of each level of DIKW.

As an intelligent system becomes capable of more complex thinking, it can convert more efficiently a given amount of emotional energy into intelligent work as compared to another intelligent system not able to think as complex. Complexity in this regard is related to curiosity, creativity and efficiency. More complex thinking systems are able to come up with (or dream up) more unique responses to stimuli.[7] This ability to be curious

---

6. This notion regarding existence suggests that there likely exists some relationship between emotion, intelligence and consciousness that manifest once an intelligent system is capable of higher level thinking (i.e., K and W). While we refrain from speculating more about this relationship, we do conclude this report with a final thought on this notion.

7. There is an argument to be made that similar to protein folding, where the entropy lost due to folding into a complex, ordered nature is overcome by increasing the entropy in its local environment (i.e. hydrophobic effect), intelligent systems, through

or creative gives the intelligent system a better chance at understanding the mysteries of its environment. This leads to a higher probability it will evolve to a more complex state that is more efficient in converting energy into work that is used to build and maintain its complex, ordered nature.

This explanation suggests that all intelligent systems are driven to more complex states of intelligence processing by an emotional force. This force provides the motivation to optimize its use of energy to maintain its complex, ordered nature. Additionally, it provides the drive to build and improve its intelligence models that explain itself and the world around it, thus increasing its ability for complex intelligence processing (i.e., intelligence complexity).

Without an emotional force to drive the intelligence process, any intelligent system will lack the motivation or the drive to develop the higher levels of thinking described by the DIKW hierarchy. This is why we propose that a fundamental relationship between intelligence (I), emotion (E), and the intelligence complexity (C) of an intelligent system must exist.

CONCEPT: Achieving Machine Emotion

The focus of this report thus far has been on understanding intelligence and emotion in a more universal context, as it applies to any system capable of intelligence and emotion. There is however, a real push in today's world for interpreting intelligence and emotion in the context of machines. Companies, like IBM and Microsoft are investing billions every year into creating machines with higher levels of intelligence and therefore, as this report would argue, it is important to consider how

---

more complex thinking, are able to increase the disorder of its local environment. In this regard, higher intelligent systems capable of more complex thinking end up leading to a greater measure of disorder locally due to the "creative and curious" changes they make in the environment. See also Alexander Wissner-Gross and Cameron Freer, "Causal Entropic Forces," *Phys Rev Lett* 110 (Apr19), no. 16 (2013): 168702.

to provide the driving force for machines to achieve higher levels of intelligence. As the previous section discussed, we think this means investigating the idea of producing emotion in a machine, just as it would for any system.

So, if we are to achieve machine intelligence at the knowledge level (or wisdom level, collectively higher level intelligence), then the machine will need the drive, in a form that we have called emotion, to build models. If we focus on the knowledge level, then the machine will need a dose of what we have called creativity.

At the information level, the emotion that fuels curiosity is much more easily programmed into a computer. In fact, Google's Deep Blue supercomputer recently was shown to demonstrate "artificial curiosity," by programming a kind of reward system for making more random choices.[8] Indeed, the curiosity can emanate from the programmer, or the user, in queries that are imposed to the machine that then inspects the database to retrieve the information. Alternatively, the program may simply say: Find trends in the data. Whether the trends are based on regression analysis or complex topological features of the data is irrelevant: It is the program that drives the extraction and assemblage of information. The extraction of information might be better if the machine itself had the "curiosity" to look for interesting trends and associations in the data, independent of its explicit programming or specific user-generated queries. But curiosity at this level would not require "consciousness," nor a full range of emotions and motivations.

There are in fact many examples of artificial intelligence systems that operate at an information level, that are able to generate information from queries and trends analysis. Many of these systems depend on being fed massive sets of data that contains many examples of any given

8. George Dvorsky, "Artificial Curiosity Allows This Bot to Triumph at *Montezuma's Revenge*" *Gizmodo*, June 7, 2016, http://gizmodo.com/artificial-curiosity-allows-this-bot-to-triumph- at-mont-1781067908.

phenomena. In this sense, the systems are information level intelligent, and are driven by the sheer volume of data. Big data analytics is based on the idea that more data is better, and that systems using big data will be able to find information that might not be readily apparent to a human. They accomplish this in part because computational speeds of machines are so much faster than humans, but more specifically because they are able to assimilate large bodies of data rapidly.

In this sense, one might argue that emotion is inherent in their architecture. Still, most AI systems must be told what information to look for, or at least how to extract information. In that sense, the machines are simply tools that are assisting humans in finding and extracting information. Thus it is debatable as to whether the machine or the human is producing intelligence at the information level. We can generally ascribe information level intelligence to these machines, simply because they are so valuable in assisting the extraction of information that a human might not find independently, and thus the machines at least seem intelligent.

Accordingly, to achieve higher-level machine intelligence, where the machine truly has the measure of intelligence at the designated level, we believe it will be necessary to provide positive drive (or emotion, as we have been calling it). This does not mean that the machines need to have the full range of human emotions, or be "sentient," or be able to pass the Turing Test. What it means is that the machine, either through its programming or architecture, has to have sufficient drive so as to be able to create models (or in the case of wisdom, meta-models).

Consciousness, Intelligence and Emotion

In order to achieve consciousness, an intelligent system likely needs to achieve a certain level of complex thinking and be driven by some motivational force like emotion. According to our theory of DIKW, the level of complex thinking required for an intelligent system to obtain a minimum level of consciousness would be at least knowledge. This is because it is only at this level of complexity that an intelligent system can construct models that describe the data and information it can process. When the system begins to construct a model about the data and information that corresponds to its "self," our theory would argue that system has achieved some level of self-awareness and thus consciousness.

Beyond this notion that consciousness is related to an ability for an intelligent system to create a model of "self" via an emotional driving force guided by complex thinking, we refrain from speculating anymore about the true nature of this relationship in this report. As with our theories regarding both DIKW and Emotion, we hope this introduction to a possible theory to explain intelligence and consciousness can act as a catalyst to spur the right conversations moving forward.

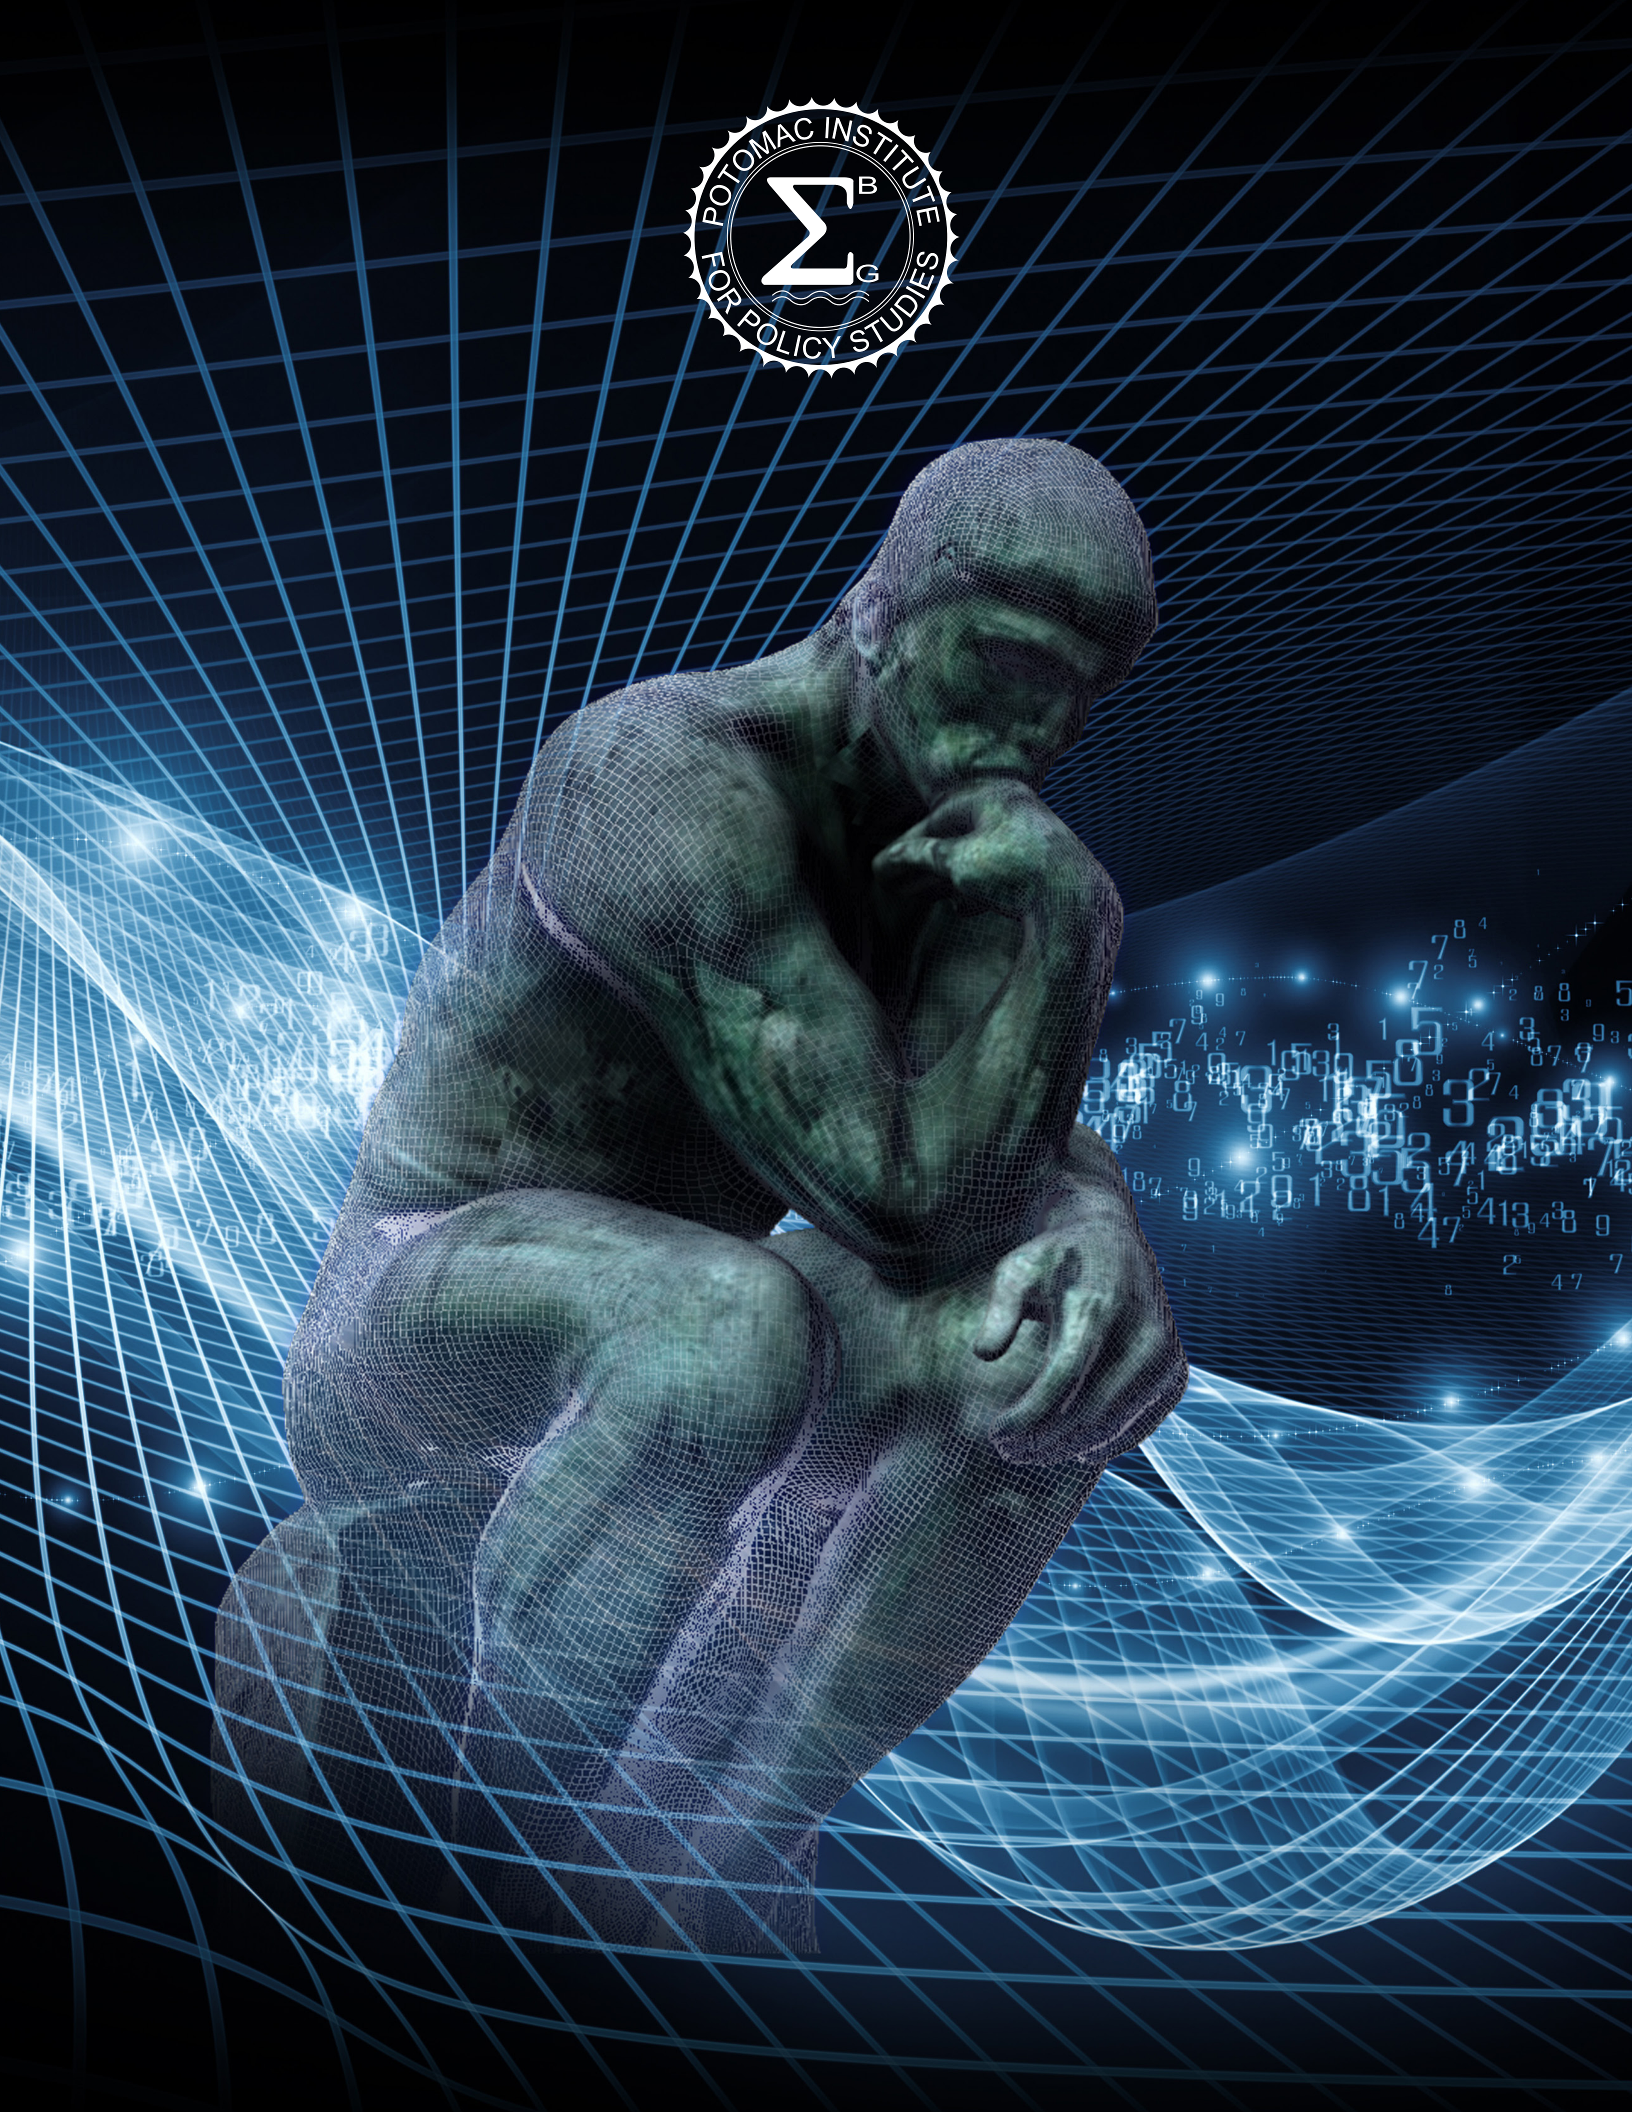
POTOMAC INSTITUTE
FOR POLICY STUDIES