


How Managers Should Prepare for Deep Learning: New Values



Exploring the unique challenges for managing analytic systems enabled by Deep Learning [photo by [Charlie Wild](#) on [Unsplash](#)]

[Richard Hackathorn](#), Aug 20, 2018

TL;DR—see Summary at end for quick overview.

This article is the second in [a series about the Managerial Perspectives on Deep Learning](#), which are targeted toward managers who are responsible for or involve with analytical systems enabled by **Deep Learning (DL)** using artificial neural network technology. They faced confusing concepts and unique challenges when dealing with these systems. This article focuses on the **New Values** that suggest different approaches for evaluating the value realized by an organization from DL-enabled analytics. At the end of each section, a Take-Away offers practical suggestions for managers.

Toward Intelligent Effectiveness

The value that large IT systems generates for an organization and society in general is shifting from **cost efficiency** to **intelligent effectiveness**.

As Peter Drucker stated, “*Efficiency is doing things right; effectiveness is doing the right things.*” [01] Leveraging this quote, cost efficiency is performing specific tasks faster, cheaper and reliably, while intelligent effectiveness is customizing tasks to match the needs of the specific situation.

Both are important but focusing of effectiveness is better when the purpose of task can not be fully specified or significantly varies with each situation. This is especially true for systems at large scale.

Since the early days of computing, the value of computer systems has primarily been on efficiency with the goal of cost reduction. Manual procedures were error-prone in data capture and reporting. An early computer application was simply to print payroll checks on paper, which was a huge efficiency improvement over manually writing checks. The task of generating a payroll was precisely defined, thus ripe for this computerization.

Today, eCommerce systems are assisting thousands of customers in selecting the right product from thousands of choices, based on wide combinations of factors, such as price, bestselling, lowest cost, recently updated,

associated purchases, and review stars. There are no pre-specified procedures since features are constantly varying dependent on product availability and customer uniqueness. Being intelligently effective is essential to generating value in this large scale.

Take-Away: As a manager involved with analytical systems, become aware of opportunities for customizing standard procedures for unique situations. Assess the potential business value, and highlight these opportunities to LOB colleagues. Ask technical colleagues to think about how analytics could be utilized so that each situation is treated uniquely, at over a thousand per hour. Foster a list of these opportunities for intelligent effectiveness within your organization.

Generalizing Beyond Known Data

The intelligent effectiveness of these systems depends on understanding ecosystem behavior on the large scale but interacting uniquely on a small scale. The enabling ingredient is **generalizing beyond known data**, which contrasts with **describing known data**. [02]

The definition of *generalizing* implies making a statement (hypothesis) that is applicable (valid) for a broad range of situations, by inferring from examples of a smaller (but representative) sampling.

This is the essence of analytics—to take data about a set of example cases and create a model applicable to similar cases. There are two key assumptions behind analytics that generalize beyond known data. First, the example cases to train the model must be representative of all the cases that the model will predict. Second, the trained model must be validated on a random (with-held) subset of examples for accurate predictions.

In normal life, humans with their intuition have been *generalizing beyond known data* since the dawn of civilization. In fact, our survival as a species has been dependent upon this ability. More recently, we have used our IT systems to describe known data in ingenious ways, using regular reports, interactive dashboards, and visual analytics. Thus, human intuition has excelled at creating insights that lead to generalizing causal relationships and predicting business futures.

Take-Away: As a manager involved with analytical systems, are you aware of areas within your organization that have exceeded the ability of human intuition to be managed properly? Can you cite specific examples of this failure? What are the factors that have led to this problem? Is your business treading into new areas for which many years of business experience are not relevant to current business problems? Is there data that can adequately describe examples for this current problem?

A Million-Unique-Things

When dealing with intelligent effectiveness situations, there are limits for human intuition to be consistent and reliable, even when enhanced with descriptive analysis tools. As the complexity of systems increases, studies have shown that intuition is inherently irrational and, hence, sub-optimal. [03]

One measure of complexity from a managerial perspective is for that manager to be responsible for a *million-unique-things* (or MUT for short). [04] Let's consider three examples to illustrate this point.

Imagine you are an owner of a small retail store. You have been doing business in your community for years and know personally most customers visiting your store. You instinctively treat these customers uniquely with customized service, while maintaining a sense of the community as a whole. When a stranger enters your store, you greet them warmly and acquire information about their unique needs. You probably get an 'A' on

intelligent effectiveness. So, how many unique customers do you have? Probably several hundred. Rounding to powers of ten, the number of unique customers is roughly around 10^2 to 10^3 for a small retail store.

Imagine you are successful and build your business to multiple larger stores through your region. Now your unique customers have increased maybe to 10^5 . Do you personally know most customers visiting your stores today? Probably not. How's is your intelligent effectiveness doing with serving your customers?

Imagine you are wildly successful and build your business to a global retail chain with thousands of stores operating in dozens of countries. Now the number is easily over 10^6 . To serve your customers, you now have a real big problem with intelligent effectiveness! You must deal with your MUT!

The conclusion is that, **around a million-unique-things, human intuition to generalize beyond known data become unreliable.** [05]

The complexity of our world has now exceeded our abilities of human intuition to generalize reliably at scale, especially in economic, social, and political areas.

Analytics (with its ability to generalize beyond known data at scale) becomes necessary. The technology to support these analytics has evolved from descriptive statistics (over the last one hundred years) to machine learning (over the last twenty years) to artificial neural networks (over the last five years).

Note that, in the big perspective of human history, this technology evolution is an extremely recent development! The ancient Greeks did not have a MUT problem.

As a human race, we must embrace wisely analytics to enable us to be intelligently effective with coping with the increasing complexity of our world.

As argued in a related article, analytics should be used to augment human intuition, as opposed to replacing it with automated processes. [augment intuition] The challenge is to augment human intuition with objective data-driven support for generalizing beyond known data.

Take-Away: As a manager involved with analytical systems, survey the critical areas of your organization that must deal with MUTs. Look for hints in the major subject areas of your integrated data warehouse. Is your organization performing adequately in these areas? If your organization could substantially improve its intelligent effectiveness in these areas, would it result in a significant benefit?

Volume—Velocity—Variety

You do not have to be a global company with a thousand stores to have the requirement of understanding and managing a MUT. The term *big data* is often misused to sell the promise of a successful business for you and a better life for all. However, the 3V's concept from big data is useful to explain why analytics can enable intelligent effectiveness. It all depends on whether adequate data exists to describe those MUTs.

The real issue is dealing with the *information content* required to manage a situation involving a MUT. [06] On other words, how many bits of data are needed to characterize adequately a MUT? One gigabyte, or one petabyte? [06]

Using the example of a million customers who shop at your stores, what does a MUT mean in terms of the data collected about these customers?

Consider the 3V's dimensions—**Volume** ÷ **Velocity** ÷ **Variety**—to explain information content of data and, thus, its potential to support generalizing beyond known data.

Volume... The example of customers dealt with the volume of things as the number of unique things—its cardinality. However, this can be tricky depending how the customer is defined. Is a customer an individual person, or a family, or a large corporation? Further, is the thing a customer interaction with the store? And, is this interaction a physical visit to a store, or a browse of the website, or a response to mail coupon, or a reply to an email offer? Hence, you may think your organization deals with a few hundred customers, but you will be forced dealt with millions of things to understand those customers.

Velocity... The example of customer interactions deals with the velocity of things—unique things per unit time. Whenever the time dimension is relevant (as it often is), then data velocity becomes an important factor. All customer relationship management (CRM) systems perform analysis on customer interactions, not customers. The results of the analysis may be aggregated to individual customers for mailing campaigns, but the results are based on generalizing these interactions.

Variety... The example of customer as a person, family, or corporation deals with the variety of the thing—information richness of a single thing. The information required to describe adequately a person, family, or corporation is widely different. If your customer is IBM or Apple, you will need lots of data to understand how to serve that corporate customer properly. Further, data is now enriched with the rich variety inherent within images, text, audio, sensors, and the like.

Therefore, the information content of a MUT is some product of volume, velocity and variety of those things. For instance, a hundred things (such as managing corporate customers, instead of individual persons) could be sufficient to exceed that IC level where human intuition hits the wall. As humans, we were not designed (nor evolved) to surpass this level. [07]

Take-Away: As a manager involved with analytical systems, how are the 3V's affecting your information content for critical million-unique-things areas?

Examples with Features (plus Labels)

From conventional machine learning to the newer neural networks, a common characteristic is the format of the input data, which is *examples* (or rows, instances) of *features* (columns, attributes), plus a *label* (or target) as one of the features. This is called the *example-feature-label* (or EFL for short) format for data. [08]

Most think of EFL data in a tabular format, given that it is often illustrated in an Excel-like spreadsheet. However, that is misleading since a feature for a DL-enabled model could be an image, music, or other unstructured data. Current implementations of neural networks process *tensors* (multi-dimensional matrices) as the basic unit of data.

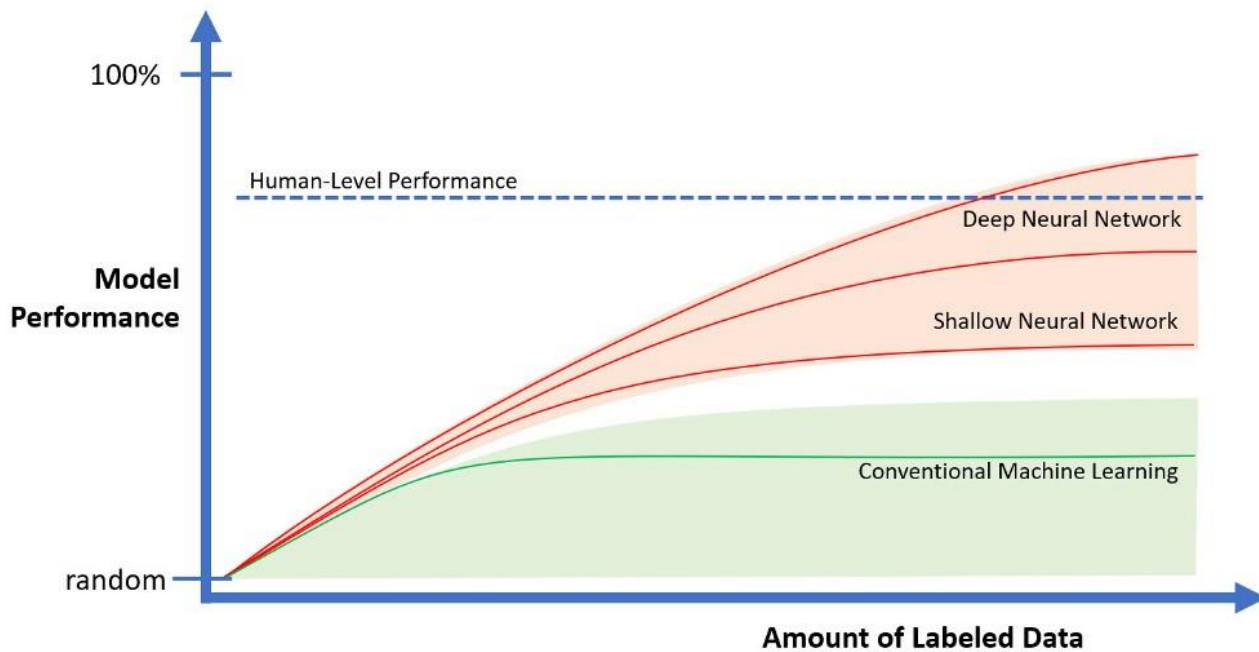
One of the features may be a *label* (or target) that indicates an important characteristic for this example. For instance, a customer would be represented by an example consisting of various features, such as age, sex, and the like. The label could be the sum of their purchases over the past month. The label feature is a target to 'supervise' (guide) the algorithms optimize toward the valid generalization.

In summary, data analytics begins (and ends) with *examples* of things-of-interest, described by a set of *features*, one which could be a *label* of an important characteristic. The label is often a categorical value, such as 'very important customer', that is usually determine by subjective judgment.

Take-Away: As a manager involved with analytical systems, concentrate on the curation of the label in EFL data. It is a crucial statement of the desired outcome of the model, given data similar to the other features. Its value is likely set subjectively by experts, which is expensive, error-prone, and scale inflexible.

EFL Data as the Value Driver

In the article on New Paradigms, this figure was explained. [09] The horizontal axis corresponds to the amount of EFL data, with the right side being the MUT area. The vertical axis corresponds to how well the analytics generate a model that can generalize beyond known data.



Therefore, this figure (in the upper right) shows that neural networks are currently the best technology to deal with the MUT problem. Further, this technology can exceed human-level performance for some tasks. However, the key requirement is that we can collect and curate sufficient EFL data about our MUTs.

Take-Away: As a manager involved with analytical systems, assess the amount of EFL data that can be curated for MUT problems. If insufficient data, will conventional machine learning techniques be adequate? If there is sufficient data, does the potential business value motivate your organization to exploit neural networks? Do you have the resources to do so?

Analytic Maturity Stages

The value of analytics is often described as a sequence of maturity stages, for which certain questions can be answered. [10] Note that the latter stages are various forms of generalizing beyond known data.

1. **Descriptive**—What happened? ...to describe known data
2. **Diagnostic**—Why did it happen? ...to infer causality from correlations
3. **Predictive**—What will happen? ...to infer aspects of future events
4. **Prescriptive**—What should happen? ...to optimize a plan toward desirable objectives

This framework implies... A company achieves greater analytical maturity by evolving its analytical capabilities from stage 1 to stage 4. Since questions of greater generalization are enabled, the greater analytical maturity increases the potential for generating greater business value. However, the effort (along with costs, skills, resources, risks) also increases.

The maturity stages are useful for evaluating analytical capabilities in terms of the types of questions that it enables. It relies heavily on human intuition derived subjectively from insights, which may not be sufficient to deal with MUT problems, as previously discussed.

Take-Away: As a manager involved with analytical systems, assess the ability of IT systems to enable the analytic maturity stages across various functional areas. Is your analytics adequate to handle MUT problems for current operations and for future growth?

Analytic Value Chain

In organizational studies, the framework called a “value chain” defines value as the sequence of activities required to deliver a product or service to the consumer. The concept was first described by [Michael Porter](#) in his 1985 book [Competitive Advantage](#) to explain the functions of a typical corporate in terms of its ability to generate value to its stockholders.

Let’s apply this paradigm to analytics in terms of an Analytic Value Chain.



Analytic Value Chain as pipeline from data to action

Analytic Value Chain is a pipeline from data to action, resulting in an increasing value in the form of intelligently effective organizational behavior. The sequence flows from left to right (with considerable looping in actual practice). The raw materials are the data, while the finished goods are actions, which are intended to change the behavior of the company.

If your organization must deal with managing MUTs, then it must be supported the capabilities of all five stages, with the end-to-end result being data-to-action. The implications are: Business value from analytics is created when they enable business processes to be intelligently effective to affect positively customers, products, and the like. If there are no positive impacts on business processes via actions, then no benefits will occur for the organization. With these inputs, **All the analytical systems are just costly overhead!**

Take-Away: As a manager involved with analytical systems, managers should adopt an end-to-end focus since this framework represents a flow of value that could be interrupted at any stage. Also, the middle three stages are the focus of the data science teams. Hence, managers should pay special attention to the first and last stages. The first—Acquire & Curate—deals with the issues discussed in previous sections, such as MUT and EFL. The last—Operationalize & Govern—are the subject of the next section.

Analytic Last-Mile

As with large electrical, communications, and delivery systems, the proverbial ‘last mile’ is often the most expensive, slowest, and error prone of any segment within those complex systems. The same is proving true for analytics when viewed as an end-to-end value chain.

The analytic last-mile emphasizes the last stage of the value chain—Operationalize & Govern—which deal with operationalizing the analytic, along with governing its application. At this point, the initial EFL data was used to train and validate a neural network model. The model enhances a specific business process in an intelligently effective fashion, and its behavior is monitored and govern according to a set of policies.

Current practice is often called *Analytic DevOps* to indicate the close synergism between the teams developing the model and those executing its production application. With neural network models, a set of unique issues are emerging:

- Can the required software execute in the production environment? What is required to do a model prediction?
- Operationalizing real-time model prediction within large streaming systems, instead of batch prediction updates to data warehouse
- Detection that the current model is no longer valid because operational data has diverges from training data
- Whether the model needs to be retrained: Can retraining be performed concurrently within the production environment? At what time cycles—daily, hourly?
- Detection of unintended bias in training data, such as sexual/racial bias.
- Linkages to data governance to manage models as learning logic, not static logic
- Continual refinement of model performance via champion-challenger dynamic
- Interpretability of model results for managers and for affected consumers
- Management and auditability of the entire end-to-end analytic workflow

Take-Away: As a manager involved with analytical systems, the analytic last-mile presents your greatest challenge. Allocate sufficient resources early and often, along with a high priority for your attention.

Learning as the Value Sustainer

The final topic is that learning by the analytical system becomes the long-term sustainer of value. Once that learning ceases, the value of that system declines because the ability of the model to generalize about new examples diminishes.

This paradigm shift is best captured by Andrej Karpathy, Director of AI for Testa, in his 2017 article [Software 2.0](#). He writes...

Because it is now easy to collect data about many business problems, neural network models can be trained using this data faster and efficiently, as compared to writing explicit static code. These models are literally data-driven, rather than code-driven as the prior 50 years of business computing.

Karpathy makes the distinction between *static logic* (hard-coded) with *learned logic* (training by examples), assuming that sufficient data about those examples exists.

Let’s take it one more step to *learning logic* with continuous re-training on new example data.

The usual situation is... A model is trained and validated on an original set of examples. Creating that example set is a difficult, costly and labor-intensive effort. After the model is operationalized into production system, it performs as expected. And, the model is NOT re-validated on new examples because of the high effort to create those new examples.

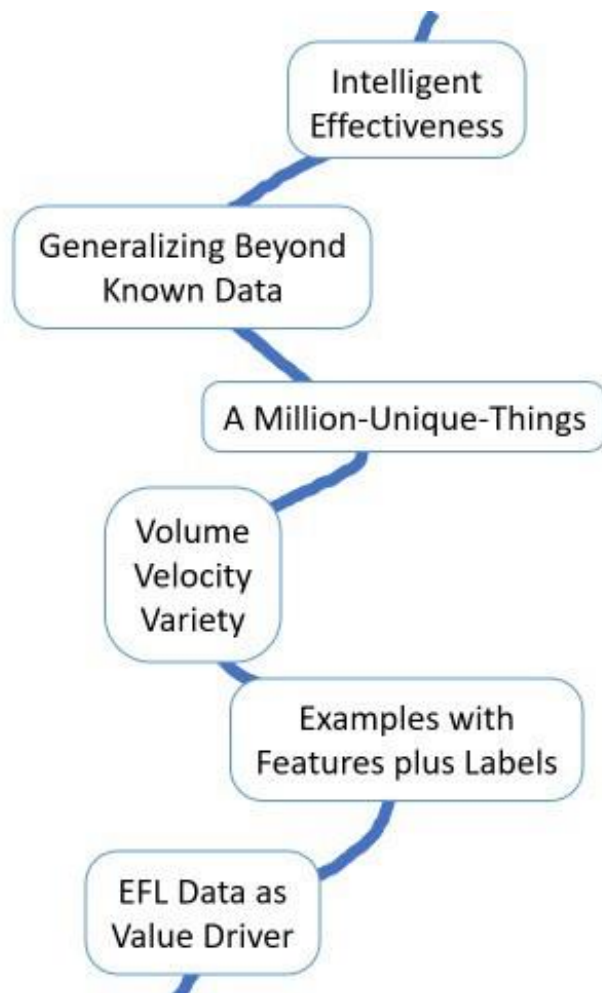
Note the underlying assumption... New examples are assumed to be similar (representative) to the older ones because the business has not changed significantly. Hence, it is not worth the effort to re-train.

However, the current environments for most organizations is changing more than is realized. And, the data infrastructures are maturing so that new examples can be collected as part of normal operations, if that capability is implemented as part of the operationalizing the model.

Take-Away: As a manager involved with analytical systems, what is the business value of moving your analytical systems from static logic to learned logic and to learning logic? Do you have the resources to realize that value?

Summary

This article builds a step-by-step argument for defining and determining the value of a large analytical system. The argument explains these steps in basic concepts that can be discussed across all management levels and provide a common basis for evaluation and collaboration.



The argument starts with contrasting effectiveness (doing the right task) with efficiency (doing the task right). The focus is set on *intelligent effectiveness*—customizing the task to match the needs of the specific situation. The critical capability is to *generalize new situations, based on data about past situations*, thus enabling customization of the tasks. Human intuition has served this function well, until recently when the complexity has exceeded its limit. This limit roughly occurs when a situation involves managing more than a *million-unique-things* (or MUT for short). Yet, things come with variations in *volume, velocity and variety* so a few complex things coming fast can be very complexity to manage.

The format for capturing this complexity are a set of *examples*, each with its *features*, plus a *label* indicating a target value. Data in this format of Example-Features-Label is called EFL data, which is the value driver of analytical systems. As the amount of EFL data increases, DL models perform increasingly well, eventually exceeding human-level-performance levels.



Discussion shifts to four *analytic maturity stages* dealing with various questions. This is contrasted with a *analytic value chain* that illustrates the value-added pipeline for transforming data into actions. Highlighted is the analytic last-mile, which is the Operationalize & Govern step as often being the bottleneck for generating value for the organization.

The final section focuses on *learning as the value sustainer*. A distinction among static logic, learned logic, and learning logic is made. The latter being the only alternative for sustaining value over the long-term.

Appreciation... to the DL study group at [Data Detectives of Boulder](#) for our discussions over past months, which helped clarify these thoughts.

Finally... If you benefit from these articles, please support [my Patreon](#) to create and mentor small peer groups of managers to explore key management issues of analytical systems enabled by Deep Learning. If this program might be of interest to a colleague, please share a link or tweet. Thanks, Richard

Endnotes

[01] The quote is from Drucker's book [The Effective Manager: Getting the Right Things Done](#). A concise [summary of the 12 lessons](#) from this book seems quite applicable to managing DL-enabled systems. Also, a no-nonsense elaboration of this Efficiency-Effectiveness distinction was captured in this May 1993 article [Managing for Business Effectiveness](#).

[02] Generalization is a fundamental concept in science, logic and statistics for hundreds of years. The [definition from Wikipedia](#) captures its essence: *formulation of general concepts from specific instances by abstracting common properties*, which leads to defining the interrelated parts that form the whole. The term abstracting is the outcome of a process that follows general rules, like that of a data analytics algorithm.

[03] The 2008 book [Predictably Irrational: The Hidden Forces That Shape Our Decisions](#) by Dan Ariely is a revolutionary book about how today's social dynamics, individually or societally, appear to behave in ways counter to rational logic. Hence, models based on rationality are misleading, at best. In 2016, [Jono Bacon of Forbes interviewed Ariely](#) to probe *how we [can] help machine to understand us [humans] better*. One quote stood out: *Rather than trying to work against human nature I think we should be working with human nature and ... look at the mistakes people make and ... think about what is the path to fix that [mistakes] and what kind of systems that would involve*.

[04] The term *thing* is intentionally informal and used to refer universally to objects or events involved in a situation. The formal term would be *entity*, which is defined in [Merriam-Webster](#) as anything with *independent and separate existence*. Within data structuring, the term entity has a formal role in the [entity-relationship model](#) (ER model) as the *things of interest* that are interrelated within a situation.

[05] **TBD: Need research** to support the assertion *humans cannot generalize reliably beyond a MUT*. Reading the 2012 book [Too Big to Know: Rethinking Knowledge](#) by David Weinberger for inspiration.

[06] The term [*information content*](#) (or self-information) is a foundational topic for analytics, involving [*information theory*](#) by Claude Shannon in 1948 and formalized as the concept of [*entropy*](#). It is the ‘surprise’ obtained from a random sample of data. The more surprise, the more entropy. If the surprise from the data doubles (i.e., offering twice the alternatives) and continues to double (e.g., times 2, 4, 8...), then the number of bits within the information (i.e., its entropy) is 1, 2, 3...

[07] The solution is to augment human intuition, as opposed to replacing it with automated processes. This point—Augment Not Automate—was explained in the [*New Paradigms article*](#). Analytics can be the tool to reduce information complexity to a level that human intuition can comprehend and perform successfully. The analogy is like a telescope or microscope that augment human senses and intuition to see and understand very large or very small objects. In note [03] above, augmentation should assist with mistakes in human intuition.

[08] This is assuming the usual *supervised* (in contrast to *unsupervised*) training of the model. The training is supervised (directed, evaluated) by the accuracy of predicting the ‘true’ label value. The set of examples is divided into a larger *training set* and a smaller *testing set*. Each cycle through all training examples is called an *epoch*. After each epoch, the mismatches in label and predicted values for both training and testing examples is noted and plotted in a *learning curve*. This chart shows how well the model is generalizing from the examples, in both *bias* (hitting around the right value) and *variance* (hitting in a tight cluster around some value) errors. This summarizes most of conventional machine learning practice today. However, deep learning practices are expanding this framework in surprising ways, as will be discussed in a future article *New Horizons*, which discusses *meta-learning* research.

[09] This chart was explained in more detail in the Data Drives Performance section of the [*New Paradigms article*](#). It is a simplified version of the slide titled “Scale Drives DL Progress” by Andrew Ng in the [*Coursera DL Specialization*](#).

[10] Most charts showing four-stages of analytic maturity are attributed to Gartner in March 2012. However, some cited an article by Tom Davenport in 2007. Even Wikipedia is silent on its origins. Perhaps the chart should be renamed to analytic capability stages (or levels).