




Ethical Use of Artificial Intelligence for Actuaries





Ethical Use of Artificial Intelligence for Actuaries

AUTHOR

Neil Raden
Founder/Principal Analyst
Hired Brains Research



Give us your feedback!
Take a short survey on this report.

[Click here](#)

**Caveat and Disclaimer**

The opinions expressed and conclusions reached by the authors are their own and do not represent any official position or opinion of the Society of Actuaries or its members. The Society of Actuaries makes no representation or warranty to the accuracy of the information.

Copyright © 2019 Society of Actuaries. All Rights Reserved.

CONTENTS

1	Executive Summary	5
2	AI, the Social Context and the Five Pillars of Ethical AI	7
2.1	Defining AI	7
2.2	The Social Context	8
2.3	The Five Pillars of Ethical AI.....	10
2.3.1	Responsibility.....	10
2.3.2	Transparency.....	11
2.3.3	Predictability.....	11
2.3.4	Auditability.....	11
2.3.5	Incorruptibility.....	11
2.4	Macro AI Ethics Obligations	11
3	Digital Transformation Drives AI Adoption	12
3.1	The Motivation for Digital Transformation.....	12
3.2	Future Technologies and InsurTech.....	13
4	Ethical Risk in General: Data, Bias and Do-It-Yourself Risk	14
4.1	Defining Ethical Risk with AI	14
4.2	Ethical Considerations for Actuaries Using AI	15
4.2.1	Accelerated Underwriting.....	15
4.2.2	Access to Third-Party Data beyond Marketing Data	15
4.2.3	Reliance on Spurious Medical Markers from New Sources	15
4.3	Data Risk.....	16
4.4	Bias.....	16
4.4.1	“People” questions to ask about AI Bias	17
4.4.2	Data Questions to Ask about AI Bias	18
4.4.3	Management Questions to Ask about AI bias	18
4.5	Ethical Aspects of Automation	18
4.6	Personally Identifiable Information (PII).....	18
4.7	Reliance on Repeatable Algorithms without Human Oversight.....	19
4.8	Amateurish Development	19
4.9	Insurance Circular Letter No. 1 (2019)	20
5	Organizing the Actuarial Department for AI: Skills, Team, Diversity.....	21
5.1	Skills Needed for Actuarial Development of AI	21
5.1.1	GLM (Generalized Linear Model)	21
5.1.2	Technology around AI	21
5.1	Technology Skills for Actuaries.....	22
5.2	Graph Analysis.....	22
5.2.1	Pricing Analysis.....	23
5.2.2	Agent Compensation	23
5.3	IT and AI Engineer Team.....	23
5.4	IT Role in Actuarial AI Team.....	23
5.5	Ethical Review Board	25
6	Government and Regulatory Initiatives Concerning Ethics in AI	26
7	Conclusion: Suggested Path Forward	27
8	Appendix 1: Advanced Analytics Categories	28
8.1	What Is Predictive Analytics?.....	28
8.2	What Is Data Science?	28

8.3	Predictive Analytics versus Data Science.....	28
8.4	What Is Machine Learning?.....	29
8.5	Some Misconceptions about Machine Learning.....	29
9	Appendix II: Examples of Review Boards.....	30
9.1	American Psychological Association.....	30
9.2	PSI (Population Services international).....	30
9.3	National Society of Professional Engineers.....	31
9.4	Swedbank.....	31
9.5	Software Developers.....	31
9.6	Ethics Review Board: Roles Required.....	31
9.7	Other Codes of Ethics Examples.....	32
10	Additional Suggestions for Reading.....	33
10.1	Light Reading.....	33
10.2	Ethics Articles by Neil Raden.....	33
10.3	More In-Depth.....	33
10.4	Books.....	34
10.5	Resources for Frameworks for Ethical Analysis.....	35
11	Endnotes.....	36

1 Executive Summary

Actuaries understand quantitative modeling. It seems only natural that they would adopt Artificial Intelligence (AI) as an extension to their already extensive training and skill in predictive and probabilistic modeling. However, in addition to requiring more learning to be effective, AI imposes requirements that are very different from current analytical modeling. AI is fueled by data—and great volumes of it in a multitude of forms. In addition, because of its ability to both operate in an unattended mode as well as mask its internal logic, AI can pose ethical risks in its implementation.

The scope of this paper is to provide a technical overview of the tools and disciplines currently in AI as well as the forces at work that financial institutions such as insurance companies are using to modernize their analytical processes. The aim is to provide some guidance to the actuarial profession for understanding ethics in relation to using AI, recognizing the potential for doing harm and how to avoid it. The purpose of this report is to highlight the ethical risks arising from the application of AI in actuarial practice and to have tools to use to identify and manage it.

AI, the Social Context and Ethics

Ethics in AI becomes an issue when the social context is involved. To limit unintended consequences arising from AI implementation, companies can incorporate guidelines into their AI development and implementation processes. Formulating and practicing the ethical application of AI in actuarial work requires consideration of some simple [guidelines](#)¹ that can be useful for formulating a more detailed policy, such as the following:

- AI researchers should acknowledge their work can be used maliciously
- Policymakers need to learn from technical experts about these threats
- The AI world must learn from cybersecurity experts about how to protect its systems
- Ethical frameworks for AI need to be developed and followed
- More people need to be involved in these discussions. It should not just be AI scientists and policymakers, but also businesses and the general public.

Some primary sources of AI risks include the following:

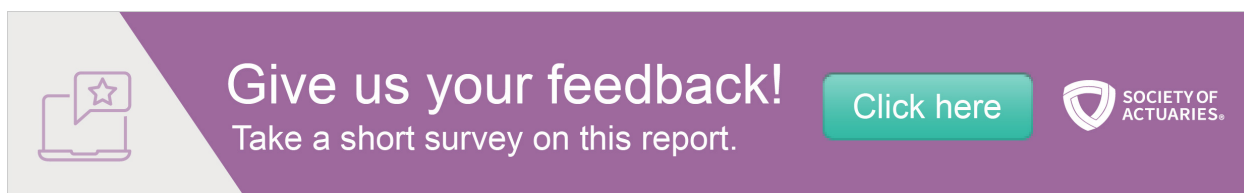
- Data risks, resulting from imperfect or misunderstood data
- Bias risks, resulting from unintentional bias that can manifest itself at many points in the AI development process
- Automation risks, where the automation of a process can have unintended consequences
- Amateurish development, where the systems are not developed by AI professionals and unintended consequences result.

To address the risks associated with an AI system implementation, this paper proposes a five-point “Five Pillars of Ethical AI” as a high-level guide:

- Creating **Responsibility** for what a developer creates and uses
- Using **Transparency** to ensure the logic of an AI is viewable
- Ensuring that output from an AI implementation has **Predictability** and produces consistent results
- Guaranteeing **Auditability** of the outcomes

- Ensuring that AI systems have **Incorruptibility** and are protected from manipulation.


These five pillars provide a fundamental oversight framework for developing and implementing an AI system. By considering any implementation from the viewpoint of each one of these pillars, risks associated with AI implementations can be minimized.



The banner features a purple background with a white diagonal stripe on the left. On the left side, there is an icon of a laptop with a star in a speech bubble. The main text is in white, and there is a green button with white text. On the right, there is the Society of Actuaries logo.

Give us your feedback!
Take a short survey on this report.

[Click here](#)

 SOCIETY OF ACTUARIES

2 AI, the Social Context and the Five Pillars of Ethical AI

Several concepts need to be defined to clarify our ethics in AI discussion. First, we present a brief description of AI, in a broad sense as it is applied to actuarial work. Then we introduce a discussion of ethics and propose a set of basic principles for the Five Pillars of Ethical AI. All of this discussion would meander without nailing it down to a context, and that is the social context: when AI meets people.

2.1 Defining AI

AI is not a new technology. The earliest principles of AI are more than 50 years old. When AI first became a commercial offering and AI start-ups appeared in the 1980s, the capacity to compute the algorithms was so severely limited that AI went into “AI Winter,” not to be seen again for 20 years. However, the rise of very large-scale analytics, hybrid cloud architectures and a general explosion of computing capacity provided for economical examination of data and models at a scale not previously possible. Interestingly, many of the algorithms used for analytical work are the same as those used in statistical models. For example, both a statistical model and machine learning (ML) model may use logistic regression. The ML model gradient methods, such as optimization, consume very many compute cycles, while classic statistical modeling is mostly sampling and mathematical equation solving. Actuaries generally have some facility for the latter, but it is less likely that they are skilled in the former.

The AI discipline is moving toward areas of cognition, such as natural language processing (NLP), image recognition and robotics. For the purposes of this paper, we will confine our analysis to those aspects of AI that are most relevant to actuarial work: AI enhanced data management/discovery, analytics, learning algorithms and supervised and unsupervised machine learning. For the sake of clarity, we call this “narrow AI.” Some practitioners may use the term “available AI,” and they typically mean the same thing. Because of industry confusion, there isn’t a precise definition of the terms AI, ML, data science and predictive analytics (see Appendix I, “Advanced Analytics Categories”).

When we define the social context shortly, it makes sense to include other digital decision-making disciplines in the framework of consideration. For the sake of inclusion, all sorts of decision-making disciplines can be bundled into “Available AI,” such as decision management (including rules and decision trees), predictive analytics, ML, data science, deep learning, NLP and natural language generation (e.g., chatbots).

The following represents a reasonably clear consensus for defining narrow AI:

- Systems and applications that perform tasks that mimic or augment human intelligence and decision-making agents that adapt to their environment
- AI is not the technologies that enable it— including ML, deep learning, neural networks, natural language processing, rule engines, robotic process automation and combinations of these capabilities for higher-level applications
- Defining “intelligence” is tricky; logic, reasoning, conceptualization, self-awareness, learning, emotional knowledge, planning, creativity, abstract thinking and problem solving are all elements. In addition, beyond those, ideas of self, of sentience and of being can also be included. Artificial Intelligence is, therefore, a machine that possesses one or many of these characteristics.

Keep in mind that most AI can do more than a human may be able. It can be faster as well but generally can consider only one thing at a time. Facial recognition software may pick out a face and match it with a name, but it has no memory or emotion about that person. The same is true with ML. It can learn only what it is taught. AI often can be thought of in categories. The holy grail of AI is **General AI**, a single system that’s equal or superior to human intelligence. Today’s AI applications do not exhibit this “humanlike” intelligence—an algorithm designed to drive a car, for example, would not be able to detect a face in a crowd or guide a domestic robot assistant or have an opinion about Flaubert’s *Madame Bovary*. The ethical considerations of General AI are far more complex than

Narrow AI. Narrow artificial intelligence is a much more specific use, where a technology outperforms humans in some very narrowly defined task. Narrow AI is more the subject of this research.

ML implementations are classified into three major categories, depending on the nature of the learning “signal” or “response” available to a learning system:

- **Supervised learning:** When an algorithm learns from example data and associated target responses to predict the correct answer when given new examples. The developer provides good examples for the algorithm to memorize and then derives general rules from these specific examples.
- **Unsupervised learning:** When an algorithm learns from clear examples without any associated response, leaving to the algorithm to determine the data patterns on its own. They are useful in providing insights into the meaning of data and new valuable inputs to supervised ML algorithms.
- **Reinforcement learning:** When an algorithm is given examples that lack labels (descriptions of the data) but can have examples with positive or negative feedback. The algorithm must make decisions (so the product is prescriptive, not just descriptive, as in unsupervised learning), and the decisions bear consequences. In the human world, it is just like learning by trial and error.

The output of ML models can be either of two classifications when inputs are divided into two or more classes:

- **Regression**, which is also a supervised problem, a case when the outputs are continuous rather than discrete or
- **Clustering**, where there is a set of inputs to be divided into groups. Unlike classification, the groups are not known beforehand, making this typically an unsupervised task.

Some additional misconceptions about ML can be found in Appendix II.

From an ethical point of view, however, all forms of advanced analytics, such as predictive models and data science, are subject to the same ethical standards as AI so long as the object of their investigations involves people, a subject described below as the “social context.”

2.2 The Social Context

Ethics in AI becomes an issue when the social context is involved. Social context² refers to the immediate physical and social setting in which people live or in which something happens or develops. It includes the culture that the individual was educated or lived in and the people and institutions with whom they interact. Almost everything an actuary does involves a social context because it involves people. As soon as AI is placed in a social context, ethical questions come into play.

An autonomous drilling machine knowing when to oil itself has no social context. The use of AI in Human Resources within organizations can be within the social context in many situations, such as an algorithm using specific hiring parameters, and not considering things that a person would.

In a recent issue of *Strategy + Business*,³ there is a thought-provoking piece entitled “[Team Human vs. Team AI](#).” The article notes:

We shape our technologies at the moment of conception, but from that point forward, they shape us. We humans designed the telephone, but from then on, the telephone influenced how we communicated, conducted business, and conceived of the world. We also invented the automobile, but then rebuilt our cities around automotive travel and our geopolitics around fossil fuels.

Artificial intelligence adds another twist. After we launch technologies related to AI and machine learning, they not only shape us, but they also begin to shape themselves. We give them an initial goal, then give them all the data they need to figure out how to accomplish it. From that point forward, we humans no

longer fully understand how an AI program may be processing information or modifying its tactics. The AI isn't conscious enough to tell us. It's just trying everything and hanging onto what works for the initial goal, regardless of its other consequences.

It's an interesting perspective because it provokes us to think about the long-term effects and the monitoring we should be exercising regarding AI-related technologies. In [How to Boost Artificial Intelligence Education in Your Company](#),⁴ Ben Lamm of *Forbes* writes, "AI still has bias problems; deep learning still gets stumped by simple, if unexpected, variables; neural networks are far from the point where they can be consistently leveraged strategically for business."

Nan DeMars, workplace ethicist and author of "You've Got to Be Kidding: How to Keep Your Job and Your Integrity,"⁵ says that in ethical businesses, employees need to feel free to ask questions and know that their concerns are taken seriously and kept confidential. They also need to feel confident that there will be no repercussions from communicating about perceived problems.

AI provides many new opportunities for ethical issues in practice beyond current practices. Among them are "black box" decision models whose internal processes are hidden and can mask all sort of bias, and new forms of unregulated data and algorithms that were not historically part of the actuarial training (though elements of this have been added more recently) and therefore not as well understood when implemented. Actuaries have always been balancing risk management and considering ethical issues. For example, if an actuary uses data and an AI model to determine risk scoring system for all properties in a city (see, e.g., <http://cytora.com/ai-guide/>), how do they ensure the AI isn't inadvertently building discriminatory rating (using unallowed rating variables)—or using combinations of variables that would be proxy discriminatory? Additional examples can also be found at <https://www.insurancenexus.com/analytics/artificial-intelligence-and-actuary-future>.

Insurance companies today can access hundreds of facts about individuals in underwriting processes for very low cost. In addition, making a good faith attempt to remove variables from a model that can be biased may still be influenced by [latent inputs](#).⁶ In other words, one cannot say "Let's just exclude credit from the model" while including 20 other variables that may help encode someone's creditworthiness. Even when these attributes are not used in a learning model, the attributes are latent values⁷ that lead the model right back to an undesirable outcome. The opportunity to misuse that information, deliberately or not, is an ethical issue of AI today. Without an ethical framework for the use and governance of AI, including oversight of data quality management, potential biases or incorrect conclusions can work insidiously into the actuary's work process.

Kathy O'Neil, in her bestseller [Weapons of Math Destruction](#),⁸ gives a comparable example:

A college student with bipolar disorder wanted a summer job bagging groceries. Every store he applied to was using the same psychometric evaluation software to screen candidates, and he was rejected from every store. This captures another danger of AI: even though humans often have similar biases, not all humans will make the same decision. However, given the same inputs, the inferencing algorithm will always make the same decision. That is desirable, but only if humans are to judge for themselves. Perhaps that college student would have been able to find one place to hire him, even if some of the people making decisions had biases about mental health.

It is interesting that in this entire space of AI and analytics, the discussion around values is supercharged. This has to do with the fact that analytics and AI are potent weapons that can be used in very strategic and very targeted ways.

Moreover, as a result of this, it seems crucial for an organization that chooses to implement these techniques to have a set of values that governs these techniques. Just because one can do something

does not mean that one actually ought to do it. With these potent tools available that can have unintended consequences, there has to be an obligation to understand the broader impact.

2.3 The Five Pillars of Ethical AI

At times, the insurance industry has [struggled to keep up with pace of new technology](#).⁹ The pressures of digital transformation—together with the promise of increased efficiency, new products and services and new markets—comes with increased risk, especially from the application of AI. A well-funded international business of [InsurTech threatens to upset practices from customer facing services to actuarial modeling](#).¹⁰ Moving from manual methods to unattended decision-making systems can subject an actuarial organization to risk from bias, privacy intrusions, security and issues of fairness. It should be clear that AI is fundamentally different from predictive models and data science; however, predictive analytics, without AI, is rapidly becoming more accepted in actuarial work and needs to be subject to ethical scrutiny as well. The dividing line between what is prudent risk evaluation and what is unethical gets fuzzier with AI.

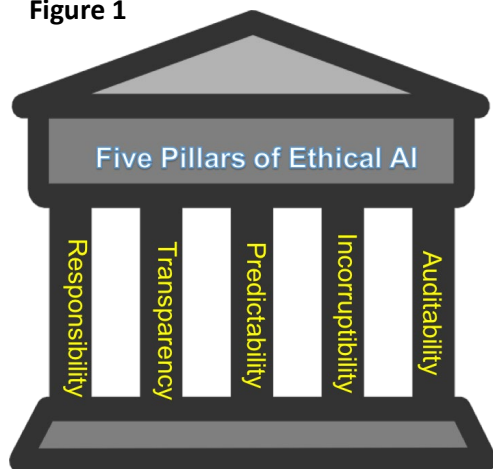
The discriminatory effect of good intentions presents the obvious ethical issue: “I can build it, but should I?”

The roles of actuaries will also change, and updated guidelines for ethical practice are often needed. A clear statement of ethical principles is needed for actuarial practice using AI. There are nascent efforts to create ethical frameworks, not always specifically for actuarial practice, but they provide a good starting point. Actuarial societies have a long history of providing helpful guidance to the practice. This paper deals with what emerging or additional guidance for AI implementation can be helpful.

At its most fundamental level, an ethical approach to the application of AI rests on five foundational concepts, shown as pillars in Figure 1.

A discussion of these five pillars appeared in [“Can AI Be Bounded by an Ethical Framework,”](#)¹¹ in September 2018, where the topic was first discussed. Quite a few similar articles have been written about “pillars” since then. The pillars are general in nature and are not specific to actuarial practice, but foundational. Further development of ethics for AI in actuarial practice should derive from these five concepts.

Figure 1



2.3.1 Responsibility

Who is responsible for the ethical behavior of the AI? In a colorfully titled article, [“When an AI Finally Kills Someone, Who Will Be Responsible,”](#)¹² the author describes the work of Gabriel Hallevy at Ono Academic College in Israel in “Artificial Intelligence and Legal Liability,” which examines whether criminal liability could ever apply, to whom it might apply and, under civil law, whether an AI program is a product that is subject to product design legislation or a service to which the tort of negligence applies. This article questions and explains three types of liability that could attach to an AI implementation. The first is criminal liability, which usually requires an action and a mental intent (in legalese, *an actus rea* and *mens rea*). The second is natural probable consequence, which occurs when the ordinary actions of an AI system might be used inappropriately to perform a criminal act. The third is direct liability, and this requires both an action and an intent.

The first pillar of ethical AI is identifying who is responsible. In the case of actuarial practice, responsibility rests with the entire team that produces an application based on AI. This even includes using embedded AI capabilities of packaged software such as data discovery/management tools or analytical algorithms.

2.3.2 Transparency

The second pillar of ethical AI is ensuring transparency. Can an action be explained? If conclusions are derived from specific techniques, such as deep learning neural nets or genetic algorithms, it is currently difficult to explain how the findings were reached. Other methods, such as decision trees and Bayesian nets, are much more transparent. Transparency should be at the core of the use of AI within actuarial practice.

2.3.3 Predictability

The third pillar of ethical AI is creating a predictable environment. Does the AI provide a predictable environment where people can optimize their own lives, or does it encourage helpless frustration for people? Unfortunate implementations of AI include examples as simple as the AI creating prompts to questions that would be out of context or help systems that require a long-forgotten password. The list of unpredictable occurrences that stem from the use of AI can be endless if not considered correctly.

2.3.4 Auditability

The fourth pillar of ethical AI is ensuring auditability. The ability to continuously track the actions of the AI through a facility that is part of AI itself (or otherwise open to inspection by third party) is an important consideration.

2.3.5 Incorruptibility

The fifth pillar of ethical AI is ensuring incorruptibility. Robustness against manipulation is a common requirement in computer security systems but less likely to be of emphasis in AI applications. Incorruptibility should be at the core of the use of AI within actuarial practice.

A scan of the literature reveals that these five principles are more or less uniformly agreed upon. Some discussions are more technical, others more fully extrapolated. IBM, for example, has issued a steady drip feed of papers and discussion documents on the ethics topic. In a recent report, “Everyday Ethics for Artificial Intelligence,” it suggests a series of principles that are not a million miles away from the five above. Then there is a launch of IBM’s AI Fairness 360 Open Source Toolkit designed to address the issue of bias.

It is difficult to assess how committed the large technology vendors are to this topic. A job search on Glassdoor using the terms *ethical*, *ethical development*, or *ethical* turns up precious little. A search for *AI developer* triggers 745 job listings. On cursory examination of the first page of Glassdoor search results, only one listing mentioned ethics.

2.4 Macro AI Ethics Obligations

The Five Pillars provide a high-level set of guidelines for ethical practices of AI. Ethical AI obligations derive from the Five Pillars as directions for implementing ethical practices. Useful guidelines include the following:

- The potential for malicious use of AI should be clearly acknowledged
- Experts in AI need to actively engage with policymakers to inform them about the risks and remedies
- Cybersecurity experts can be invaluable in teaching AI experts how to protect systems
- The audience for these discussions should expand to the general public in addition to AI scientists, businesses and policymakers.

3 Digital Transformation Drives AI Adoption

In 2017 the U.S. insurance industry recorded \$1.2 trillion in written premiums.¹³ The premiums are split 52%/48% between Life and Annuity companies and Property/Casualty companies. Health insurers are not included in these totals. At \$1.2 trillion, insurance contributes about \$602.7 billion, or about 3.1%, to the U.S. gross domestic product (GDP). To put this in perspective, this places insurance just below the federal government contribution to GDP and just above the information technology sector. The industry as a whole has cash and invested assets over \$5.5 trillion and employs 2.5 million people. The U.S. insurance business represents about a third of the worldwide total, split more or less evenly with Europe and Asia.

Insurance is a complex industry, with many layers, disciplines, markets and regulators. It is, however, more reliant on data and analytics than most industries, with a few exceptions, but has at times struggled to keep up with new waves of technology. Of all of the sectors highlighted in case studies and sales pitches for big data analytics and AI/ML, insurance is infrequently cited if not conspicuous by its absence. Not only is it desperately in need of digital transformation, but it can also benefit from a wide assortment of technologies and applications.

3.1 The Motivation for Digital Transformation

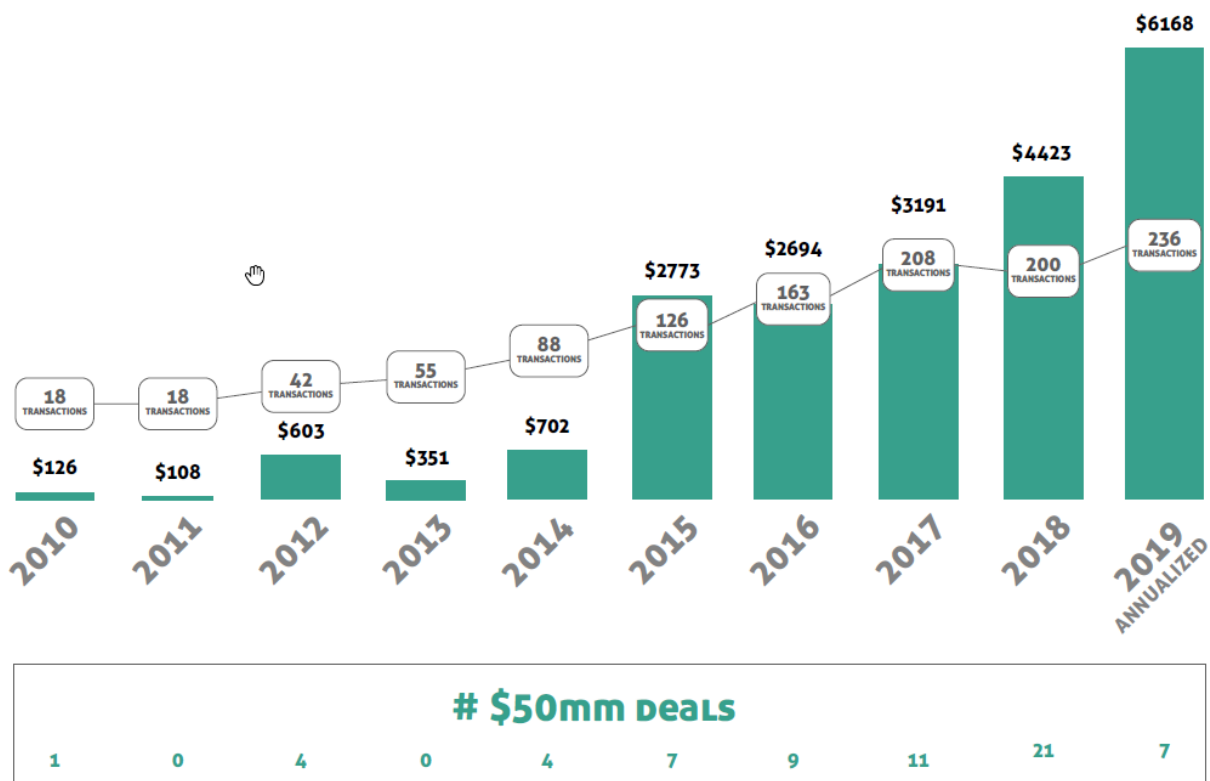
AI is an essential aspect of digital transformation of the insurance industry. Therefore, it is essential to understand what motivates organizations to invest substantial resource in digital transformation. Digital transformation is not always a precise term. One could say that all organizations have undergone a continuous digital transformation for decades. Today the primary motivations for any organization to invest substantial sums and effort in digital transformation are the following:

- **Cost:** Streamlining operations by replacing human effort with AI. Insurers look to simplify and accelerate the overall process, such as claims management procedures and underwriting.
- **Consistency:** Underwriting and claims, for example, can be completed with uniform consistency and have results that can be auditable from transaction log analysis. Underwriting and claims engage a social context, however, and social context engagement necessarily involves ethical questions.
- **Quality:** Any work that is subject to audit can conceivably be improved with better processes and AI.
- **Fraud:** ML algorithms can spot patterns of fraud that would require a great deal of effort for humans. The advantage is not only to detect fraud, but also to act on it proactively.
- **Sustainability:** Sustainable Development Goals (SDGs) “[are] the process of maintaining change in a stable environment, in which the exploitation of resources, the direction of investments, the orientation of technological development and institutional change are all in harmony and enhance both current and future potential to meet human needs and aspirations.”¹⁴ Ethics is an integral part of sustainability.¹⁵
- **Streamlining:** Interactions with insurance industry stakeholders are consistently reviewed to be streamlined for interactions with insureds, distributors, regulators and others.

3.2 Future Technologies and InsurTech

The insurance industry has often struggled to keep up with emerging technology. The pressures of digital transformation, with the promise of increased efficiency, new products and services, and new markets, comes with increased risk, especially from the application of AI. Based on the growth of InsurTech, the pressures for increased digital transformation will only increase in the future (see Figure 2).

Figure 2
InsurTech Private Financing Volume, \$ in millions



4 Ethical Risk in General: Data, Bias and Do-It-Yourself Risk

For actuaries, ethical risks arising from AI (and any advanced analytics: see Appendix I, Advanced Analytics Categories) can arise systemically through practice if not considered appropriately. Data used for ML can be biased, incomplete, not representative or poorly assembled, among other causes. Bias can also arise from assumptions of the developer or from interpretation or inference of the results. Bias can be gradual and subtle, and generally its effects can be detected only by those it harms, which makes management of its outcomes more difficult. Algorithms are not biased, people and data are, and all data should be considered within context. When data are first captured, encoding is in a context that may not always flow through its copies and transformations.

4.1 Defining Ethical Risk with AI

To the extent that these risks involve people (the social context), we can say that they are ethical risks. In general, ethics refers to standards of morally good or bad, right or wrong conduct. Defining and detecting unethical behavior or processes with AI requires attention to numerous social contexts, including:

- Discrimination
- Diversity and inclusion
- Gender equity
- Improper use of health records
- Invasion of privacy
- Lack of sustainability
- Missing governance
- Conflicts of interest
- Insensitivity
- Inequality

Life Insurance (and health insurance) are societal goods, but still in [the United States, LIMRA notes that the population is \\$15.3 trillion underinsured](#).¹⁶ Finer-grained underwriting has the potential of exacerbating the problem, not addressing it. Discussions are ongoing regarding the strength of the evidence supporting a causal role of some social factors in determining population health, such as the social determinants listed in the box. According to the Centers for Disease Control (CDC),¹⁷ “Conditions in the places where people live, learn, work, and play affect a

Table 1
Social Determinants of Health

1. Income and Social Status
2. Employment and Working Conditions
3. Education and Literacy
4. Childhood Experiences.
5. Physical Environments
6. Social Supports and Coping Skills
7. Healthy Behaviors
8. Access to Health Services
9. Biology and Genetic Endowment
10. Gender
11. Culture

wide range of health risks and outcomes. These conditions are known as social determinants of health (SDOH).” SDOH (see Table 1) are widely used in research into public health, sociology and many other fields, but for the purposes of managing risk in a life insurance company, they are too high-level to be free of significant bias. For example, why are people homeless? Perhaps they are mentally ill or have addiction problems. Or perhaps they are escaping violence or simply have found themselves in an untenable financial condition. But for the purposes of SDOH analysis, they are just “homeless.” It is a category that is too aggregated, and in quantitative studies it should not be used as a surrogate endpoint.

Meanwhile, researchers increasingly are calling into question the appropriateness of traditional criteria for assessing the evidence.

AI ethics is a much broader discussion than the origins, effects and biases in AI, and the potentially disastrous results of unmethodical experiments. Actual use and deployment in actuarial departments have the potential for ethical dilemmas.

4.2 Ethical Considerations for Actuaries Using AI

The work of actuaries and the data and tools they use are in a continuous process of change. Understanding the risks, ethical and otherwise, that come with new data sources, new modeling techniques and a new industry, InsurTech, is needed. For example, in the past, underwriting relied on a few observable criteria and judgment based on the experience of pooled classes. The ability to underwrite individuals has the potential for bias (unintentional for the most part) and for denying coverage to those who need it based on spurious data and assumptions.

4.2.1 Accelerated Underwriting

Accelerated underwriting processes, as the name implies, makes the life insurance underwriting process faster and more accessible. Typically an insurance company can issue a fully underwritten, regularly priced term life insurance policy of [up to \\$1 million without a medical exam](#).¹⁸ Accelerated underwriting is not the same as Simplified Issue, which means that there is no requirement for a physical exam, but the premium will be higher than someone who went through full medical underwriting. With accelerated underwriting, the applicant typically will have no significant medical conditions (such as cancer, diabetes or hypertension) and no biological parent or sibling who died from heart disease or cancer before age 60. Normal blood pressure and cholesterol levels are also vital.

Certain other non-health factors also are important in the accelerated underwriting process. Life companies may require no history of bankruptcy in the last five to 10 years, no history of driving recklessly or under the influence within five years, no more than two moving violations in the past three years and no felony charges or convictions. The ethical risks involved for actuaries developing accelerated underwriting processes are twofold:

1. How reliable are the data for the “nonmedical” data?
2. Are the models predicting experience reliably, and do they contain bias?

4.2.2 Access to Third-Party Data beyond Marketing Data

A predictive model or set of models taking publicly available information as inputs and providing a mortality assessment as output should lead to individuals with similar characteristics having similar risk scores. Often, data will be gathered from places such as personal demographics/sociographic¹⁹ sources, electronic health records (EHRs), credit reports, psychographic profiling and digital phenotyping. From a mortality perspective, there is almost no overlap with medical factors.²⁰

4.2.3 Reliance on Spurious Medical Markers from New Sources

Another consideration around the explosion of data and third-party models is the tendency to accept biological markers as significant when, in time, they prove to be just the opposite. Actuaries should be aware of these situations and consider their reliability. Consider prostate screening with a PSA test. The [recommendation for routine prostate cancer screening](#) from the U.S. Preventive Services Task Force (USPSTF)²¹ was recently raised from a “D” (which discourages the service since “There is moderate or high certainty that the service has no net benefit or that the harms outweigh the benefits”) to a “C” (which means that physicians should “Offer or provide this service for selected patients depending on individual circumstances” and that “There is at least moderate certainty that the net benefit is small”).

Medicine is fluid. What is considered “standard of care today” may be reversed tomorrow. [From JAMA Internal Medicine](#),²² consider: “The reversal of medical practice is not uncommon in high-impact literature: 13% of articles that make a claim about a medical practice constituted reversal in our review of 1 year of the *New England Journal of Medicine*. The range of reversals we encountered is broad and encompasses many arenas of medical practice including screening tests and all types of therapeutics.” With 400 nonmedical data points available, many from unregulated data aggregators, one has to question how reliable the assumptions in the models can be.

[Another study by the Mayo Clinic reports:](#)²³ “We reviewed 2044 original articles, 1344 of which concerned a medical practice. Of these, 981 articles (73.0%) examined a new medical practice, whereas 363 (27.0%) tested an established practice. A total of 756 articles addressing a medical practice constituted replacement, 165 were back to the drawing board, 146 were medical reversals, 138 were reaffirmations, and 139 were inconclusive. Of the 363 articles testing standard of care, 146 (40.2%) reversed that practice, whereas 138 (38.0%) reaffirmed it.”

4.3 Data Risk

Large-scale analytical tools process far more data, in many forms, for an individual to understand and vet without digital assistance. Data are too vast and too complicated to be managed by simply exporting data into databases and doing calculations in spreadsheets. They require advanced, AI-driven tools. Data are often unusable without adequate curation, governance and integration. With these in place, data can be used reliably by tools with intelligence, such as AI and knowledge graphs.

One guiding principle about data should never be overlooked: data should never be accepted on faith. How data are construed, recorded and collected is the result of human decisions about what to measure, when and where and by what methods. In fact, the context of data—why the data were collected, how they were collected and how they were transformed—is always relevant. There is no such thing as context-free data; data cannot manifest the kind of perfect objectivity that is sometimes imagined.

Geoffrey Bowker, an informatics professor at the University of California, Irvine, said, “[Raw data is both an oxymoron and a bad idea](#); to the contrary, data should be cooked with care.”²⁴

Lisa Gitelman and Virginia Jackson point out in the introduction to the aptly titled “Raw Data Is an Oxymoron” (2013), “At a certain level the collection and management of data may be said to presuppose interpretation.” “Raw data” is not merely a practical impossibility, owing to the reality of preprocessing; instead, it is a conceptual impossibility, for data collection itself already is a form of processing.

Actuaries need to be very careful about the provenance of the data they use. The term “data sharing” has, until recently, referred to scientific and academic institutions sharing data from scholarly research. The brokering or selling of information is an established industry and doesn’t fit this definition of “sharing,” but it is popping up. Scholarly data sharing is mostly free of controversy, but all other forms of so-called sharing present some concerns.

Currently, data brokers are required by federal law to maintain the privacy of a person’s data if the data are used for credit, employment, insurance or housing. Unfortunately, enforcement can be challenging, and beyond these categories there are often not extensive regulations in some countries. Although medical privacy laws prohibit doctors from sharing patient information, medical information may be available from other sources, such as from the purchase of over-the-counter drugs and pharmacy logs of prescriptions.

4.4 Bias

Bias enters models through the selection of data and by the background of the developers. AI systems developed in past decades may have contained many prejudicial biases, such as for gender and race, and would be unacceptable today. One growing counterweight to this bias is the use of broader and more diverse groups in development and testing. In some cases, actuarial groups exhibit a higher level of diversity than corresponding AI engineers. ML depends on extensive collections of data, effectively becoming the codification of history. Within these histories, there are the potential for prejudices and biases that have existed in the past.

As aspects of AI already exist in every kind of business, companies have exposure to new risks, such as bias in the AI application, fear of job loss due to automation, privacy violations, discrimination and, like any other technology,

implementations with less than ideal construction. In the article [“This Is How AI Bias Really Happens—and Why It’s So Hard to Fix,”](#) in *MIT Technology Review*, Karen Hao writes:²⁵

How computer scientists are taught to frame problems often is not compatible with the best way to think about social problems. For example, in a new paper, Andrew Selbst, a postdoc at the Data & Society Research Institute, identifies what he calls the ‘portability trap.’ Within computer science, it is considered good practice to design a system that can be used for different tasks in different contexts. ‘But what that does is ignore much social context,’ says Selbst. ‘You cannot have a system designed in Utah and then applied in Kentucky directly because different communities have different versions of fairness. Alternatively, you can’t have a system that you apply for ‘fair’ criminal justice results then applied to employment. How we think about fairness in those contexts is just totally different.

Bias can result from many sources. Some of these are discussed below.

Unintentional bias: A study of dermatological film samples of lesions was “trained” to recognize potential cancerous growths. The experiment yielded excessive false positives. Radiologists call for biopsies when lesion dimensions exceed 3 cm and use a simple 12-inch ruler. The algorithm misinterpreted the ruler for malignancy.

Data bias or intentional bias: [ProPublica analyzed COMPAS](#)²⁶ assessments for more than 7,000 arrestees in Broward County, Florida, and found the algorithm was biased against African Americans. It was two times as likely to flag black defendants as future criminals, while white defendants were labeled as low risk more often than black defendants. Only 20% of the people predicted to commit violent crimes went on to do so. The most likely cause was a problem of “feature engineering,” how important variables were selected. Whether bias is intentional or not, it is the outcome that matters.

Implicit bias: Bias flavors language; and in many cases our common language reinforces bias. Sometimes we see the writing as “a farmer and his wife are out in the field.” Both, however, are selling the crops and gathering eggs. One person does not run farms. They are both farmers. This is how implicit bias affects models in subtle ways that most people do not notice.

Insidious bias: Querying bulk data can lead to a breach of privacy. ML tools can be used to create malicious deep fakes. The CRISPR gene editing tool, which could lead to advances in genetics and agriculture, may lead to cancer or the zombie apocalypse: “The DNA damage found in a new study included deletions of thousands of DNA bases, including at spots far from the edit. Some of the deletions can silence genes that should be active and activate genes that should be silent, including cancer-causing genes.”²⁷

4.4.1 “People” questions to ask about AI Bias²⁸

One approach to reduce bias in model development is to ask questions about the people involved in building the AI system.

1. **Who is building the algorithms?** Does my team embody enough diversity in skills, background and approach? If this is something the team is lacking, it is always best to bring in more team members with different experiences who can help represent a more balanced, comprehensive approach.
2. **Do your AI & ML teams take responsibility for how their work will be used?** Do they have personal ownership over the short- and long-term impact of their technical work?
3. **Who should lead an organization’s effort to identify bias in its AI systems?** It requires people who are actively seeking to find and eliminate it.

4. How are my training data constructed? Much of the AI training data used today are constructed by humans or has some human involvement.

4.4.2 Data Questions to Ask about AI Bias²⁹

By asking certain questions about the data being used by the system, one can limit the bias that is introduced.

1. Is the data set comprehensive? Sophie Searcy, Senior Data Scientist at Metis, claims that 85% of all AI projects during the next several years will deliver flawed outcomes because of initial bias in data and algorithms.³⁰
2. Do you have multiple sources of data? You can reduce bias by incorporating multiple sources of data—or maybe just the opposite. Using multiple sources of data may reduce bias, but care is still needed to prevent the multiple sources potentially causing unintended bias.

4.4.3 Management Questions to Ask about AI bias³¹

At a higher level, management can ask question about the program to help alleviate sources of bias.

1. **What is an appropriate proportion of resources for an organization to devote to assessing potential bias?** A thoughtful organization will scale the resources dedicated to assessing bias based on the potential impact of any bias and the sensitivity to the inherent bias of the team charged with developing the AI system.
2. **Has the development team thought deeply about what metrics to use to evaluate the results?** AI systems are less likely to go off the rails if the team is evaluating them using long-term holistic metrics, including pairing metrics so that they balance each other out.
3. **How can we test for bias in training data?** One approach is to have several experts summarize the same document and then compare their results. If the result is an excellent summarization framework, these summaries will be mostly interchangeable.

4.5 Ethical Aspects of Automation

There is considerable concern that AI will cause substantial loss of jobs; clearly this is an ethical issue to be addressed. However, our research shows that many tasks make up a job, and automation can eliminate some of the dull, repetitive tasks, but it does not generally eliminate entire jobs. Most jobs, even those we think of as relatively low-end jobs, are much more complicated than we realize. Actuaries will be especially affected by the elimination of data management, spreadsheet management and even the task of writing code for specialized models that will be available as prewritten algorithms.

4.6 Personally Identifiable Information (PII)

PII is a significant cause of ethical errors in AI. Any good data scientist can merge a few different files and potentially create a data set that contains PII (see Figure 3).

Figure 3
Personally Identifiable Information

Personal Information	PII (Personally Identifiable Information)	
Name	Office location	Two types of sensitive data in sources: Personal Information and information that can be used to unmask personal data, emergent PII
SS#	Business Tel	
Drivers License	Email address	
Birth Date	Badge number	
		This data is not identifiable by merely looking at column names or other available metadata Only by looking at the data itself can an algorithm decide the data is within the realm of "sensitive" Regulatory issues: GDPR, HIPAA and more: UK/DPA, CA/CPA, ePR (EU) Organizational promises to customers and suppliers to be good stewards of data you collect about them

For example, the government releases medical records for research purposes with PII removed, but you can match people based on other databases such as the transaction records of a large pharmacy. People might be identified by a combination of prescriptions and the city. The result is that Individual people can be targeted by a model without their permission or knowledge of why they are targeted. Consider a case where a health insurer, with good intentions, communicates with members about taking their diabetes medicine, when the member had never communicated that they had diabetes. Protecting PII is not black and white: the federal [General Service Administration](#) recognizes that the definition of PII "is not anchored to any single category of information or technology. Rather it requires a case-by-case assessment of the specific risk that an individual can be identified." The broadness of this definition increases the difficulty of maintaining compliance with the PII regulations.

4.7 Reliance on Repeatable Algorithms without Human Oversight

Repeatability of an algorithm can work against people. For the most part, algorithms in production have no posterior evaluation in place; this results from the misconception that algorithms are accurate and will not make mistakes. However, algorithms "fire" at a much higher cadence than people and repeat a bias at scale (part of the appeal of algorithms is how cheap they are to use). Because users of algorithmically driven outputs do not ordinarily understand probabilities or confidence intervals (even if these are provided because the internal workings of the model are not clear), they may be comfortable accepting the algorithm's output (if they even have this option). The obvious solution is to understand how to develop better, less biased decision-making tools by combining the capabilities of humans and AI jointly.

4.8 Amateurish Development

A prominent industry analyst writes:³² "AI has been democratized. Until recently, it required a data scientist to write code. Today, with these business intelligence systems, I can point and click at a few data points, choose the variable I want to predict—like a customer's propensity to buy—and these predictive models are going to be automatically generated."

This is irresponsible and wrong. Professional statisticians can be introspective and self-critical. The American Statistical Association makes a distinction between genuine professionals and those in other fields who "pick up" from textbooks without really understanding the subject. The demand for data scientists provided an opportunity for self-identified data scientists with deficient skills to operate. Expanding their role to unsupervised ML or Deep

Learning bears a substantial ethical risk for misguided decisions based on ML's *capability to infer human behaviors, moods and beliefs (incorrectly) from available data.*

The remedy to this problem is as follows:

- Proactively build oversight and accountability controls. Construct procedures for the development, maintenance and oversight of AI upfront so errors can be detected before they occur.
- Managing AI risk and ethical concerns requires a focused approach, starting with understanding the role of AI in corporate strategy and taking action to drive value-aligned AI innovation.
- Building AI innovation will mean embedding values into the development life cycle, from use case definition, modeling and training to deployment and monitoring. Policies addressing matters such as when to put a human in the loop and when and how to conduct AI impact assessments will help prevent and mitigate any AI error.
- As much as possible, build diversity into AI teams. This does not merely mean race or origin, but rather, backgrounds, points of view and experience.

4.9 Insurance Circular Letter No. 1 (2019)

The ethical risks associated with AI are getting heightened regulatory scrutiny. The New York Department of Financial Services issued Insurance Circular Letter No. 1 (2019), on January 18, 2019, entitled "Use of External Consumer Data and Information Sources in Underwriting for Life Insurance."³³ The letter goes into great detail about the risk and attendant elements of unfairness from nontraditional data use for underwriting. Our Five Pillars are general and high-level, but the New York letter is a clear example of the expected application of responsibility, transparency, predictability, auditability and incorruptibility. It is good to keep in mind the simple adage, "I can build it, but should I?"

The focus of the letter is the use of external data "not directly related to the medical condition of the applicant that is used—in whole or in part—to supplement traditional medical underwriting." This includes the use of "lifestyle indicators" but does not include "MIB Group, Inc. member information exchange service, a motor vehicle report, or a criminal history search."

What is significant about this is that the New York Department jumped on the ethical issues of the rapidly expanding use of nontraditional data and its implications. While expressing support for the positive benefits of better and faster underwriting, the letter also cautions that "many external data sources are companies that are not subject to regulatory oversight and consumer protections, which raises significant concerns about the potential negative impact on consumers, insurers and the life insurance marketplace."

The letter then becomes very specific and places the risk precisely with the actuary and underwriter, an ethical risk in the social context in full bloom: "An insurer should not use an external data source, algorithm or predictive model for underwriting or rating purposes unless the insurer can establish that the data source does not use and is not based in any way on race, color, creed, national origin, status as a victim of domestic violence, past lawful travel, or sexual orientation in any manner, or any other protected class."

Based on its investigation, the department has determined that insurers' use of external data sources in underwriting has the strong potential to mask the forms of discrimination prohibited by these laws.

5 Organizing the Actuarial Department for AI: Skills, Team, Diversity

Actuaries have a broad knowledge of the intricacies of the business and the development of mathematical and statistical models. These skills translate well for expanding into AI, and the skills for using or developing AI applications have been more recently added to undergraduate training, exam syllabi and professional development opportunities. Some areas where actuaries may look to further increase their skill sets are the following:

1. Data management, at least the review and stewardship of AI-driven “data wrangling” platforms, especially when dealing with unstructured and external data sources
2. Ability to design ML models using AI workbenches
3. Translators: Explaining the design and results of AI models to others
4. Facility in understanding the operation of embedded AI routines in software

Actuaries should continue to be aware of professional standards and codes of conduct. An ethical question might also be “how much knowledge in AI do I need to know that I’m qualified to use AI in my work?”

5.1 Skills Needed for Actuarial Development of AI

At a minimum, beginning work with AI requires understanding the nature and operation of basic algorithms, as well as some statistical knowledge to aid in understanding and interpreting the results of these algorithms. Some of the most common algorithms currently in use include the following:

- Simple and Generalized Linear Regression
- Logistic Regression
- K-Nearest Neighbor Classification
- Support Vector Machine
- Decision Trees
- Artificial Neural Networks
- K-Means Clustering
- Naive Bayes Classification
- Recurrent Neural Networks

5.1.1 GLM (Generalized Linear Model)

Another tool, not considered AI or machine learning, is the GLM, which is widely used in many types of General Insurance pricing, but its application in actuarial work in Life Insurance, particularly complex ALM models, can be useful to peer inside AI models whose inner workings are not transparent. GLM is quite transparent; the output and consequences are much clearer. Any time an actuary ventures into using a more complex model like a neural network or other “black box” where it’s impossible to fully understand all the interactions, there should be a heightened awareness of ethical risk.

5.1.2 Technology around AI

AI development should be undertaken in a tool bench environment rather than each investigation done manually. Some tools for this are Google’s TensorFlow and offerings from DataRobot, H2O, Domino, Dataiku and many others. Knowledge of coding language is essential. Favorites come and go, but Python and R are currently reliable choices.

How extensive the skill set needs to be depends on the extent of the AI application and whether the actuaries will be responsible for the technical effort. If used as an experimental tool, to build models to understand the data, introductory AI skills will suffice but actuaries will still need deep domain knowledge and some data-wrangling

expertise. If actuaries intend to take their models and embed them in operational systems, it should be part of the [DataOps](#)³⁴ methodology.

5.1 Technology Skills for Actuaries

Insurance is a competitive business, and the relentless march of digital technology ensures that all organizations will likely adopt AI in some form to stay competitive. The accumulative effect of Moore's Law, an exponential phenomenon, pushed computing to a tipping point. The rapid transformation of computing of all kinds, but especially for analytics and prediction, put to rest the prevailing model of managing from scarcity. Instead, actuaries now have a capability to manage and analyze data unimaginable only a decade ago. The effects of this are not trivial.

Though it is difficult to generalize, actuarial practice in insurance companies has not always progressed at the same pace as digital transformations in other industries. A combination of factors, a sort of perfect storm, is bearing down on the insurance industry that will have substantial effects on actuarial practice: This includes InsurTech, Big Data, Accelerated Underwriting, Edge Computing, Hybrid Cloud Computing and an explosion of DIY (Do-It-Yourself) AI. Rather than utilizing modern analytical workflows, such as intelligent data management tools and highly functional analytical tools offering frictionless continuous intelligence from ingest to final results, some actuarial departments are still dependent on personal tools such as Excel, Access and, to a certain extent, data visualization tools like Tableau. There is, of course, a proliferation of data warehouses and data marts, but they are designed to provide curated data in scale and diversity unacceptable to ML. Although actuaries professionally manage these tools, in some cases they are inefficient and limiting. More importantly, they lack the capability and promise that current environments of predictive analytics, data science and AI provide. Of course, the introduction of these disciplines will come with critical ethical issues of bias, privacy, security and transparency.

Developing quality experiments with AI is a team effort. It is not necessary for every actuary to be fluent in every skill, but it is imperative that the team does in aggregate to avoid ethical issues.

5.2 Graph Analysis

Graphs are essential because it is impossible to navigate through the ocean of unlike data available for modeling and analysis without some tools to illuminate the process. Graph analyses are about data relationships and provide the ability to traverse faraway relations easily and quickly, something for which relational databases are quite limited. Technical metadata like column names are useful, but investigating the actual (instance) data is an important step to eliminate the risk of bias in the data, personally identifiable information and other hazards. Without robust relationship analysis in extensive collections of data, error and bias are unavoidable.

Four reasons an Enterprise Knowledge Graph Can Help You (adapted from "[What Is a Knowledge Graph](#)"³⁵) discussed the following advantages of using knowledge graphs:

1. Combine Disparate Data Silos
2. Bring Together Structured and Unstructured Data
3. Make Better Decisions by Finding Things Faster
4. Future Proof Your Database with Standards

Pricing analysis and agent compensation are two good examples of the use of graph analysis in insurance.

5.2.1 Pricing Analysis

Graph analysis can include rich metadata and attributes, not just connections between nodes. For example, in property underwriting, metadata about the weather, time, location and even connected events in time can facilitate dynamic pricing models that are more accurate forecasts rather than simple supply and demand. Graph analysis can be useful to actuaries investigating potential associations leading to fraud.

5.2.2 Agent Compensation

Another practical application of graph analysis is in financial services and life insurance, where

[shared compensation between insurance agents or financial service company representatives can be interpreted](#) as a mathematical graph with points (nodes) representing agents and the lines (links) between the agents representing shared compensation ... to provide additional input variables to predictive models that provide insight into agent behavioral aspects that other variables may not capture. Overlap ratios can be calculated directly and input variables derived from summary statistics between agents and the other agents with whom they shared compensation.³⁶

This can lead to second-degree connections for an agent which can be summarized to provide variables that provide additional behavior insights.

5.3 IT and AI Engineer Team

Business applications are programmable. In these instances, a developer enters the logic and rules explicitly. With ML, however, no direct human intervention occurs at every step, begging the question “What are the design principles?” The AI engineer must train an algorithm to have a fundamental understanding of right and wrong in a wide range of contexts. It is often not sufficient to just put the chatbot into the public sphere and say, “Here, just go learn,” without understanding the implications of how that system learns, and the subsequent consequences. Microsoft learned this the hard way with Tay, an artificial intelligence chatterbot³⁷ Microsoft Corporation originally released via Twitter on March 23, 2016; it caused subsequent controversy when the bot began to post inflammatory and offensive tweets caused by “trolls” who “attacked” the service as the bot made replies based on its interactions with people on Twitter.

Just like data scientists, many actuaries are not professional software engineers. Data scientists, in many instances may choose to proceed in a go-it-alone manner, created a backlog of unmanageable code, when they should have focused on their skill and ability in building models. Information technology (IT) departments can provide the support and infrastructure of professional services to bookend the effort with coders, infrastructure, security and implementation of the models. For actuaries to apply AI effectively to their work requires a team of software engineers, data engineers, translators and specialists with domain knowledge. Another alternative, of course, is for actuaries to leave a traditional actuarial role and broaden into a pure AI development role in the organization for other subject areas. In the former case, actuaries continue in their current role but act as team members to AI application development to support their role. In the latter, of course, they will develop horizontal skills to apply their expertise to AI efforts organization-wide.

5.4 IT Role in Actuarial AI Team

AI requires collaboration and iteration by groups with business, operational and technical skills, with a goal of business transformation. Companies look to IT to initiate and manage AI projects. AI is not like database, middleware and enterprise applications. Enterprise software evolves through well-defined, managed steps and narrow functional roles, to ensure operational continuity. IT has two critical roles that bookend AI:

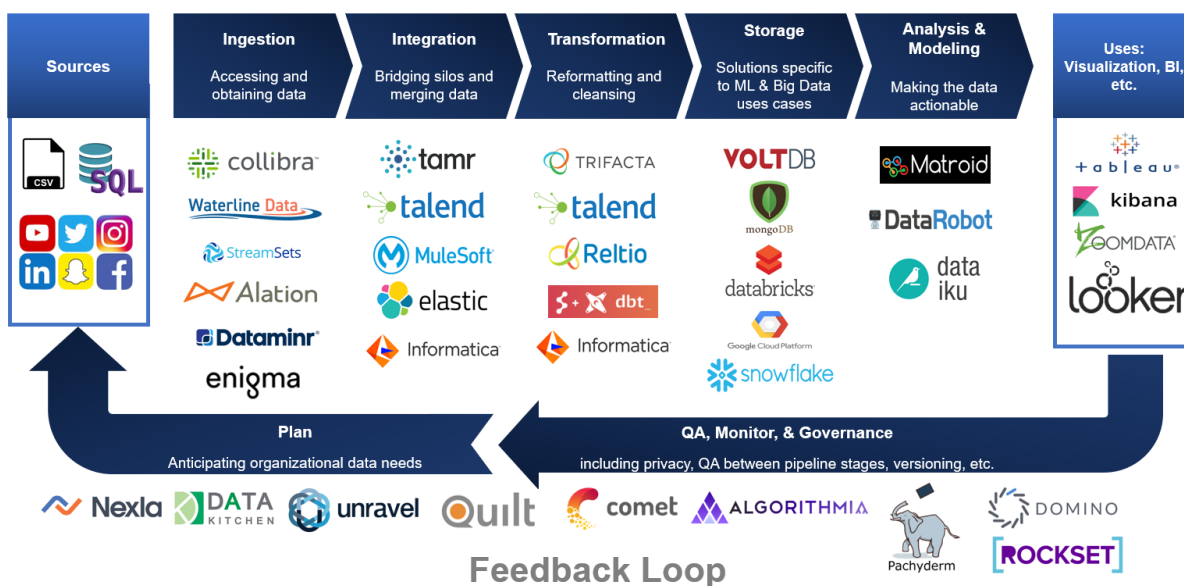
- IT database skills to acquire, integrate and govern the information assets needed for ML³⁸

- Skills for effective, reliable deployment of new AI applications.

This is, of course, a narrow definition of IT because many IT organizations have assumed the role of AI development and management. In that case, it is essential that their professional organizations develop ethical practices for AI, especially in the social context, which extends beyond actuarial practice.

Most importantly, IT has to provide the data at scale and in-time for actuaries' AI needs. The current term for this is DataOps, and it looks something like the graphic in Figure 4.

Figure 4
Data Delivery Pipeline: From Sources to Uses



Businesses need a data engineer first to organize the data, then a general data scientist who can test and adapt algorithms, and last, a software engineer who can implement the applications. Depending on the company's needs, there may be more than one of each of these positions. Collecting the right data, building the right models and using these to make the right business decisions are all important functions. More AI teams are adding project managers, marketers and User Experience/User Interface staff to round out the team. The idea here is that although data scientists and engineers have the technical skill sets, the other positions, in this case, actuaries, fully understand the product.

A truly robust and fully fleshed out AI team should be composed of the following:

- A translator or actuary communicating with the business owner to identify and prioritize the business request
- A data scientist working with the translator to develop a specific use case
- A data engineer working with the relevant business division to understand the data characteristics (requirements, source etc.) of the use case.
- A data scientist or actuary programming the algorithm and analyzing the data in the sandbox to generate insights
- A visualization analyst developing reports and dashboards for business users
- A workflow integrator working with the business owner to develop a prototype for models and tools
- A delivery manager piloting the prototype and dashboard and working to obtain a pass or fail decision.
- A delivery manager and workflow integrator working with the IT department to scale the prototype to the enterprise level.

- A delivery manager and translator or actuary working with the business and IT to ensure adoption and ongoing model maintenance.

5.5 Ethical Review Board

The formation of an organizational or professional Ethical Review Board is one approach to establishing ethical practices. The following points are suggestions for forming a board within an organization.

Based on the research of internal ethics reviews (not always specifically for AI), some universal concepts emerge for an Ethical Review Board:

- Needs the power to review any application that has a social context
- Be charged with recommending, or even acting upon, violations of the ethical code
- Composed of professionals from various departments
- Be part of the review process of any new software that may contain elements of AI
- Must moderate its work so as not to be oppressive to the progress of AI applications
- The ethical code cannot come from thin air; it has to be rooted in first principles

6 Government and Regulatory Initiatives Concerning Ethics in AI

In 2019, we have seen an expanding body of governmental and regulatory programs relating to AI ethics, either directly, or aimed at aspects that affect AI development and deployment.

European Union

In April 2019 the European Commission published a set of seven AI ethics guidelines³⁹ on how companies and governments should develop ethical applications of artificial intelligence. These principles and guidelines are similar to our Five Pillars:

1. **Human agency⁴⁰ and oversight:** AI systems should enable equitable societies by supporting human agency and fundamental rights and not be designed in such a way to decrease, limit or misguide human autonomy.
2. **Robustness and safety:** Underlying algorithms must be secure, reliable and robust enough to deal with errors or inconsistencies during all life cycle phases of AI systems.
3. **Privacy and data governance:** People have full control over their data, and the data will not be used to harm or discriminate against them.
4. **Transparency:** Visibility of AI systems is paramount.
5. **Diversity, nondiscrimination and fairness:** AI must be for everyone, regardless of color, creed, skills, abilities etc.
6. **Societal and environmental wellbeing:** AI should be used to drive positive social change and enhance sustainability and ecological responsibility.
7. **Accountability:** Mechanisms and processes should be in place to ensure human responsibility and accountability for AI systems and their outcomes.

The original report was followed by a [new report⁴¹](#) from the Commission’s independent AI High Level Expert Group (HLEG). This new report includes 33 recommendations for specific policies for European Union member states and organizations in the private sector as well as public sector bodies. The bulk of the report—Policy and Investment Recommendations for Trustworthy AI—is built around differentiating Europe’s approach to AI policy from those of the United States and China, a so-called third way that would require non-EU parties to comply.

The next phase of the HLEG’s work is to run a pilot phase of its ethical guidelines to establish their sectoral suitability and to instigate what it calls “a limited number of sectoral AI ecosystem analyses.” These should deliver more real-world conclusions than what to date has been largely theoretical work. While leaning toward the HLEG’s call for more action, much is in flux, and it’s going to be expedient not rushing into policies that are going to have a very long-term impact too hastily.

Some other regulations affecting actuaries in the United States include the following:

EU General Data Protection Regulation (GDPR): Aims primarily to give control to individuals over their personal data and to simplify the regulatory environment for international business by unifying the regulations within the EU.

NY Circular Letter: New York Department of Financial Services issued a “Circular Letter” January 18, 2019, entitled “Use of External Consumer Data and Information Sources in Underwriting for Life Insurance.” The letter goes into great detail about the risk and attendant elements of unfairness from nontraditional data use for underwriting (see the earlier discussion of this letter).

Health Insurance Portability and Accountability Act of 1996 (HIPAA): New sources of data about individuals may lead to uses that are potentially in violation of HIPAA. The danger and ethical risk is that an insurer may unknowingly include these data and reveal the data to individuals.

7 Conclusion: Suggested Path Forward

This report has discussed numerous ethical considerations that any company should think through before implementing an AI-based system or systems. We suggest below some specific points that a company should review regarding any future AI implementations.

AI systems should be unbiased:

- Consideration should be given to whether the data ingested are representative of the affected population
- Consideration should be given to whether decision-making processes introduces bias
- Significant decisions informed by the use of AI should be fair
- AI operator organizations should consider whether their AI systems are accessible and usable in a fair manner across user groups
- Consideration should be given to the effect of diversity on the development and deployment processes

AI systems should be accountable:

- Accountability for the outcomes of an AI system should not lie with the system itself
- Positive efforts should be made to identify and mitigate any significant risks inherent in the AI systems designed
- AI systems informing critical decisions should be subject to appropriate external audit
- AI subjects should be able to challenge significant automated decisions concerning them and, where appropriate, be able to opt out of such decisions
- AI systems informing significant decisions should not attempt to make value judgments on peoples' behalf without prior consent
- Diverse teams with appropriate backgrounds should develop AI systems informing significant decisions
- AI operator organizations should understand the AI systems they use sufficiently to assess their suitability for the use case and to ensure accountability and transparency

AI systems should be transparent:

- Traceability should be considered for significant decisions, especially those that have the potential to result in loss, harm or damage in the social context
- People should be informed as to the extent of their interaction with AI systems

AI systems should be as explainable as technically possible:

- AI operator organizations could consider providing affected AI subjects with a high-level explanation of how their AI system works
- AI operator organizations should consider providing affected AI subjects with a means to request explanations for specific significant decisions to the extent possible given the state of present research and the choice of model
- In the case that such explanations are available, they should be easily and quickly accessible, free of charge and user-friendly



Give us your feedback!
Take a short survey on this report.

[Click here](#)



8 Appendix 1: Advanced Analytics Categories

Advanced Analytics, Big Data and Data Science are all on an analytics “spectrum” with “narrow AI.” There are fundamental characteristics of each, but from an ethical practical perspective, they all share concerns and best practices. The ethical line between data science and AI is pretty thin. There is no unanimity about these terms; many analysts consider Machine Learning as a subset of data science. For our purposes, we do not because data scientists have more skills at data management and computer science than those who apply aspects of AI and ML to their work, whereas AI practitioners have in-depth domain knowledge relevant to their investigations and somewhat less computer science background. In either case, the results of these models are predictions that can affect the organization, people, the environment, even government policy. All of these outcomes require ethical analysis

8.1 What Is Predictive Analytics?

- Predictive Analytics involves extracting data from existing data sets to identify trends and patterns. These trends and patterns are then used to predict future outcomes and trends.
- Decision management systems treat decisions as reusable assets and introduce technology at decision points to automate the decision-making process and employ tools such as business rules, business intelligence (BI), continuous improvement (kaizen), AI and predictive analytics.

8.2 What Is Data Science?

- Initially, the term applied to high-octane statisticians and data engineers
- In practice, it turned out to bifurcate into (1) geniuses working at Google and other Silicon Valley giants and (2) data engineers who spend 80% of their time managing data, not doing “science.”
- The difference between traditional statistical, probabilistic and stochastic modeling and ML is mostly in the computation.
- Some examples of data science investigations are
 - High-frequency trading
 - Matching a Google ad with a user and a web page to maximize chances of conversion
 - Returning highly relevant results to any Google search
 - Book and friend recommendations on Amazon or Facebook
 - Tax fraud detection and detection of terrorism
 - Scoring all credit card transactions (fraud detection)
 - Computational chemistry to simulate new molecules for cancer treatment
 - Early detection of a pandemic
 - Analyzing NASA pictures to find new planets or asteroids
 - Weather forecasts

8.3 Predictive Analytics versus Data Science

- Predictive analytics is a form of quantitative modeling with existing data, expert assumptions (domain knowledge) and a wide variety of existing algorithms.
- Analysis of existing context to predict unknown events.
- Data Science employs multiple approaches to data such as data mining, data storing, data purging, data archival and data transformation.

8.4 What Is Machine Learning?

- Machine learning implementations are classified into three major categories, depending on the nature of the learning “signal” or “response” available to a learning system:
 - **Supervised learning:** When an algorithm learns from example data and associated target responses so as to predict the correct answer when given new examples. The developer provides good examples for the algorithm to memorize and then derives general rules from these specific examples.
 - **Unsupervised learning:** The algorithm learns from clear examples without any associated response, leaving the algorithm to determine the data patterns on its own. They are useful in providing people with insights into the meaning of data and new valuable inputs to supervised machine learning algorithms.
 - **Reinforcement learning:** The algorithm is given examples that lack labels (descriptions of the data) but can have examples with positive or negative feedback. The algorithm must make decisions (so the product is prescriptive, not just descriptive, as in unsupervised learning), and the decisions bear consequences. In the human world, it is just like learning by trial and error.
- The output of ML models can be either of two classifications when inputs are divided into two or more classes:
 - Regression, which is also a supervised problem, a case when the outputs are continuous rather than discrete, or
 - Clustering, where there is a set of inputs to be divided into groups. Unlike classification, the groups are not known beforehand, making this typically an unsupervised task.

8.5 Some Misconceptions about Machine Learning

- **ML Is Just about Summarizing Data**
Not True. The primary purpose of ML is to predict the future. Knowing the movies someone watched is only a means for figuring out which ones they would like to watch next. A credit report is a guide to whether someone will pay their bills on time. Learning algorithms formulate hypotheses, refine them and believe them only when their predictions come true. Learning algorithms are not yet as smart as scientists. However, they are a million times faster.
- **Learning Algorithms Just Discover Correlations between Pairs of Events**
Not true. An increase in Google searches for “flu “is an early sign it is spreading. Most learning algorithms discover much richer forms of knowledge, for example, “If a mole has irregular shape and color and is growing, then it may be skin cancer.”
- **ML cannot predict previously unseen events, a.k.a. “black swans.”**
Sort of. If something has never happened before, its predicted probability might be zero—what else could it be? However, ML can predict rare events with high accuracy. A “black swan”⁴² like the 2008 housing crisis was widely predicted – just not by the flawed risk models most banks were using.
- **The more data one has, the more likely you are to hallucinate patterns.**
True. ML experts should keep mining more attributes of the same entities to a minimum. On the other hand, mining more entities with the same set of attributes decreases the risk because the rules learned will have stronger support. For example, a person videotaping New York City Hall may not be suspicious and another buying a large quantity of ammonium nitrate may not be either, but if the two are in close phone contact, perhaps the FBI should take a look.
- **ML ignores preexisting knowledge**
Not true. The popular concept that ML is a blank slate and just crunches through data mindlessly is

incorrect. Done well, ML is the result of a long process of reasoning and experimentation. Some ML is used to refine a preexisting body of knowledge.

- **The models computers learn are incomprehensible to humans.**

If a learning algorithm is a black box, how can one trust its recommendations? Some types of models are hard to impossible to understand, like deep learning neural networks (though current research is starting to unwind this problem). However, most are easy to understand but hard to explain to nonexperts.

- **Simpler models are more accurate.**

Mostly not true. However, sometimes, the simplest hypothesis consistent with the data is less accurate for production than a more complicated one. A significant error in some ML modeling is starting with the data. There are two ways of looking at a problem: data-driven or model-driven. In the latter case, you start with what you know and build a model around it and see if the data support it. This can be simple if-then-else rules or more thorough approaches like Bayesian networks. The other alternative is to recognize some sources of data and run essentially curve fitting (most ML is curve-fitting) to see what emerges. In our consulting experience we find that being purely data-driven is more likely to introduce other effects like false positives.

9 Appendix II: Examples of Review Boards

Many professional societies and organizations have Professional Codes of Conduct. Here are four that have established and well-developed codes. These are not insurance companies and their codes do not yet address AI, but their basic principles and structure are useful to study.

9.1 American Psychological Association

The Ethics Committee shall have the power to receive, initiate and investigate complaints of unethical conduct of members (to include fellows), associate members and affiliates:

- To report on types of cases investigated with a specific description of severe or recalcitrant cases
- To dismiss or recommend action on ethical cases investigated
- To resolve cases by agreement where appropriate
- To formulate rules or principles of ethics for adoption by the association
- To formulate rules and procedures governing the conduct of the ethics or disciplinary process for approval by the Board of Directors acting on behalf of Council
- To interpret, apply and otherwise administer those rules and procedures.

9.2 PSI (Population Services international)

For its first 15 years, PSI (originally title Population Services International, but since their mission expanded after their founding they were renamed PSI) worked mostly in family planning. In 1985, it started promoting *oral rehydration therapy*. PSI's first *HIV prevention* project in 1988. PSI added *malaria* and *safe water* to its portfolio in the 1990s and *tuberculosis* in 2004. From their ethics board document:

The PSI Research Ethics Program is designed to provide “guidance to help PSI staff determine whether research activities require review by the PSI’s Research Ethics Program (REP). Considerations include whether “human subjects” are involved and, if so, whether the study is “research” according to the definitions of the USG’s Office of Human Research Protection (OHRP). Examples of PSI studies that require review as well as those that may not

require review are provided, as well as a brief summary regarding who at PSI should make these determinations.” Notice the use of the term “human subjects,” which corresponds well with our concept of “social context.”

The Research Ethics Program developed uses the following guidelines to help the staff determine when activities are required to receive ethical review and guidance on questions:

- [Our guidelines for deciding whether an activity requires REB review](#)
- [Is my activity research? A decision tree](#)
- [Request for determination of research or nonresearch](#)
- [Questions to answer before developing a submission for ethical review](#)
- <https://www.psi.org/about/>

9.3 National Society of Professional Engineers

The Board of Ethical Review is [a panel of engineering ethics experts](#) that has served as the profession’s guide through ethical dilemmas. The board consists of seven licensed members who are appointed by the NSPE president. The purpose of the BER is to render impartial opinions about the interpretation of the [NSPE Code of Ethics](#), develop materials and conduct studies relating to ethics of the engineering profession.

9.4 Swedbank

Swedbank’s Business Ethics Committee applies to the whole group, all domestic markets, and all business areas. Its role is to guide the organization to minimize business ethics and sustainability risks and any negative impacts for the bank. Anyone in the organization can escalate a case to the Business Ethics Committee. The members represent the bank’s various business areas and Group Functions. Minutes of its meetings are distributed to the CEO and the Group Executive Committee.

9.5 Software Developers

HireVue, a provider of AI-driven video interviewing and assessment tools, created a board to guide the ethical development and use of AI in its products. The board includes experts in industrial and organizational psychology, algorithmic bias, data privacy and diversity and inclusion. The company’s AI is programmed to notice factors in video interviews like facial expressions, what candidates say and how they say it.

Ultimate Software also has created a code of ethics for its use of AI: “The key is to not simply let technology ‘lead’ and deal with the ramifications later, as so often happens when we have new capabilities. In HR, we are dealing with people’s livelihoods. We need to be proactive as an industry about where we want AI to go and conversely where it should not.”

9.6 Ethics Review Board: Roles Required

Ethics Chair/Coordinator: The chair will organize, facilitate and schedule meetings; oversee the committee and ensure proper protocol is followed and appropriate interests are represented; manage the logistics of committee meetings and the deliberative process (i.e., sending out relevant case information, approving meeting minutes, documenting official decisions, etc.).

Legal Counsel: Legal counsel participation will ensure that the committee is aware of relevant authorities and limitations on those authorities based on statutes, rules and precedents. Ideally, this counsel should be local.

Local Ethics Adviser: This local ethics adviser will help provide an ethics perspective. An ethicist can be identified through either a local hospital or academic institution.

9.7 Other Codes of Ethics Examples

- The Public Health Leadership Society Principles of the Ethical Practice of Public Health (often referred to as the Public Health Code of Ethics) can serve as an essential resource
- Ethical analysis frameworks provide a process for identifying, analyzing and resolving ethical and values dilemmas. Several frameworks are available.
- Ethical analysis framework from Ruth Gaare Bernheim, Phillip Nieburg and Richard J. Bonnie:
- Identify relevant codes of ethics; In addition to the Public Health Code of Ethics, some examples of other codes one might want to draw upon including the following:
- American College of Epidemiology Ethics Guidelines
- American Dental Association Code of Ethics American Medical Association Code of Ethics
- American Nurses Association Code of Ethics Belmont Report (ethical principles for the protection of human research subjects)
- National Commission for Health Education Credentialing Code of Ethics
- National Environmental Health Association Code of Ethics

10 Additional Suggestions for Reading

10.1 Light Reading

Kozyrkov, Cassie. Are you using the term “AI” correctly? *Hackernoon.com*, May 26, 2018, <https://hackernoon.com/are-you-using-the-term-ai-incorrectly-911ac23ab4f5>.

Kozyrkov, Cassie. Explainable AI won’t deliver. Here’s why. *Hackernoon.com*, November 16, 2018, <https://hackernoon.com/explainable-ai-wont-deliver-here-s-why-6738f54216be>.

What is AI bias? <https://hackernoon.com/tagged/what-is-ai-bias>.

Kozyrkov, Cassie. Why businesses fail at machine learning. *Hackernoon.com*, June 28, 2018, <https://hackernoon.com/why-businesses-fail-at-machine-learning-fbff41c4d5db>.

10.2 Ethics Articles by Neil Raden

Howlett, Den. Can AI be bounded by an ethical framework? *Diginomica.com*, September 23, 2018, <https://diginomica.com/can-ai-be-bounded-by-an-ethical-framework/>.

Raden, Neil. Retrofitting AI—key adoption issues in the enterprise 2019–2020. *Diginomica.com*, January 2, 2019, <https://diginomica.com/the-retrofitting-outlook-for-ai-in-the-enterprise-2019-2020/>.

Raden, Neil. How hard is it to solve the ethics in AI problem? *Diginomica.com*, February 11, 2019, <https://diginomica.com/how-hard-is-it-to-solve-the-ethics-in-ai-problem/>.

Raden, Neil. Getting closer to guidelines on ethical AI. *Diginomica.com*, February 20, 2019, <https://diginomica.com/getting-closer-to-guidelines-on-ethical-ai/>.

Raden, Neil. Thinking about thinking. *Diginomica.com*, March 27, 2019, <https://diginomica.com/thinking-about-thinking/>.

Raden, Neil. Data for good—a question of relevancy. *Diginomica.com*, April 1, 2019, <https://diginomica.com/data-for-good-a-question-of-relevancy/>.

10.3 More In-Depth

Vincent, James. The problem with AI ethics. *The Verge*, April 3, 2019, <https://www.theverge.com/2019/4/3/18293410/ai-artificial-intelligence-ethics-boards-charters-problem-big-tech>. See especially the comments.

Wagner, Ben. Ethics as an escape from regulation: from ethics-washing to ethics-shopping? *PrivacyLab.at*, 2019, https://www.privacylab.at/wp-content/uploads/2018/07/Ben_Wagner_Ethics-as-an-Escape-from-Regulation_2018_BW9.pdf.

Hao, Karen. This is how AI bias really happens—and why it’s so hard to fix. *MIT Technology Review*, February 4, 2019, <https://www.technologyreview.com/s/612876/this-is-how-ai-bias-really-happensand-why-its-so-hard-to-fix/>.

Floridi, Luciano. AI4People’s ethical framework for a good society: opportunities, risks, principles, and recommendations. *eismd.eu*, <http://www.eismd.eu/wp-content/uploads/2019/03/AI4People%E2%80%99s-Ethical-Framework-for-a-Good-AI-Society.pdf>.

Johnson, Bobbie, and Gideon Lichfield. Hey Google, sorry you lost your ethics council, so we made one for you. *MIT Technology Review*, April 6, 2019, <https://www.technologyreview.com/s/613281/google-cancels-ateac-ai-ethics-council-what-next/?set=>.

West, Darrell M. The role of corporations in addressing AI’s ethical dilemmas. *Brookings.edu*, September 13, 2018, <https://www.brookings.edu/research/how-to-address-ai-ethical-dilemmas/>.

Beckers, Sander. AAAI: an argument against artificial intelligence. *AAAI.org*, 2017, <https://aaai.org/ocs/index.php/WS/AAAIW17/paper/view/15084/14649>.

Goldsmith, Judy, and Emanuelle Burton. Why teaching ethics to AI practitioners is important. *AAAI.org*, 2017, <https://aaai.org/ocs/index.php/WS/AAAIW17/paper/view/15210/14652>.

Global Pulse and iapp. Building ethics into privacy frameworks for big data and AI. https://iapp.org/media/pdf/resource_center/BUILDING-ETHICS-INTO-PRIVACY-FRAMEWORKS-FOR-BIG-DATA-AND-AI-UN-Global-Pulse-IAPP.pdf.

Editorial Board. The “deep fake” threat: high-tech forged videos could wreak havoc on politics. Policy makers must be ready. *Bloomberg.com*, June 13, 2018, <https://www.bloomberg.com/opinion/articles/2018-06-13/the-deep-fake-video-threat>.

Gralla, Preston. Agricultural genomics: feeding a growing, hungry world. *HPE.com*, November 14, 2017, <https://www.hpe.com/us/en/insights/articles/agricultural-genomics-feeding-a-growing-hungry-world-1711.html>.

Ethics guidelines for trustworthy AI. [file:///Users/mac/Downloads/AIEthicsGuidelinespdf%20\(1\).pdf](file:///Users/mac/Downloads/AIEthicsGuidelinespdf%20(1).pdf).

10.4 Books

Broussard, Meredith. 2018. *Artificial Unintelligence: How Computers Misunderstand the World*. Cambridge, MA: MIT Press.

DeMars, Nan. 2011. *You’ve Got to Be Kidding: How to Keep Your Job and Your Integrity*. Hoboken, NJ: Wiley.

Gitelman, Lisa. 2013. *Raw Data Is an Oxymoron*. Cambridge, MA: MIT Press.

Mirchandini, Vinnie. 2010. *The New Polymath*. Hoboken, NJ: Wiley.

Noble, Safiya Umoja. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press.

O’Neil, Cathy. 2016. *Weapons of Math Destruction*. New York: Crown.

Pearl, Judea. 2018. *The Book of Why: The New Science of Cause and Effects*. New York: Basic Books.

Taylor, James. 2012. *Decision Management Systems: A Practical Guide to Using Business Rules and Predictive Analytics*. IBM Press/eParson.

Vallor, Shannon. 2016. *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting*. New York: Oxford University Press.

10.5 Resources for Frameworks for Ethical Analysis

Choose an ethical analysis format that works for you. Some possible resources include:

Edited by Anna C. Mastroianni, Jeffrey P. Kahn, and Nancy E. Kass, *Ethics and the Practice of Public Health* (Oxford University Press), 2019.

Bulger, Ruth, *The Ethical Dimensions of the Biological and Health Sciences* (Northwest Association for Biomedical Research) (pp. 349–352) , Cambridge University Press, 2002.

[A Framework for Ethical Decision Making](https://www.scu.edu/ethics/ethics-resources/ethical-decision-making/a-framework-for-ethical-decision-making/), Markkula Center for Applied Ethics at Santa Clara University.
<https://www.scu.edu/ethics/ethics-resources/ethical-decision-making/a-framework-for-ethical-decision-making/>

11 Endnotes

¹ <https://www.theverge.com/2018/2/20/17032228/ai-artificial-intelligence-threat-report-malicious-uses>

² https://en.wikipedia.org/wiki/Social_environment

³ <https://www.strategy-business.com/article/Team-Human-vs-Team-AI?gko=f1c4c>

⁴ <https://www.forbes.com/sites/forbesdallascouncil/2019/07/17/how-to-boost-artificial-intelligence-education-in-your-company/#1e0f0b866718>

⁶ In statistics, latent variables, are variables that are not directly observed but are instead inferred from other variables that are followed. Mathematical models that aim to explain observed variables in terms of latent variables are called [latent variable models](#).

⁷ At the expense of over-simplification, latent features are “hidden” features to distinguish them from observed features. Latent features are computed from observed features using matrix factorization. An example would be text document analysis. “Words” extracted from the documents are features. If you factorize the data of words you can find “topics,” where “topic” is a group of words with semantic relevance. Low-rank matrix factorization maps several rows (observed features) to a smaller set of rows (latent features). To elaborate, the document could have observed features (words) like [sail-boat, schooner, yacht, steamer, cruiser] which would “factorize” to a latent feature (topic) like “ship” and “boat.”

⁸ [*Weapons of Math Destruction*](#) (New York: Crown, 2016).

⁹ <https://www.information-age.com/insurance-sector-falling-behind-others-comes-technology-expectations-123464563/>

¹⁰ <https://www.cnbc.com/2017/09/13/insurance-industry-faces-disruption-from-insurtech-report.html>

¹¹ <https://diginomica.com/can-ai-be-bounded-by-an-ethical-framework>

¹² <https://arxiv.org/abs/1802.07782>

¹³ <https://www.iii.org/fact-statistic/facts-statistics-industry-overview>

¹⁴ “Sustainability as an ethical/moral principle requires that decisions should be made that not only focus on current actions but also anticipate potential, expectable consequences”; <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4934536/>

¹⁵ <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4934536/>

¹⁶ <https://www.iii.org/fact-statistic/facts-statistics-industry-overview>

¹⁷ <https://www.cdc.gov/socialdeterminants/index.htm>

¹⁸ <https://www.investopedia.com/insurance/accelerated-underwriting-easy-life-insurance/>

¹⁹ Sociographics take the target market down to the level of the individual, where they determine the specific values, attitudes, friends, hobbies, passions and influences. Essentially, they allow you to discover what really pushes your customers’ buttons. <http://www.fionamceachran.com/understanding-sociographics-in-market-research/>

²⁰ Minyu Cao, FSA, and Phillip A. Adams, FSA, Session 16: Credit-Based Scores: Why They Work and Their Consequences, 2019 Life and Annuity Symposium.

²¹ Is PSA now “OK”? What the task force said about the evidence on prostate cancer screening; <https://www.healthnewsreview.org/2017/04/new-prostate-cancer-screening-guidelines-are-an-update-not-a-capitulation-as-portrayed-in-some-news-reports/>

²² <https://jamanetwork.com/journals/jamainternalmedicine/fullarticle/1105961>

²³ [https://www.mayoclinicproceedings.org/article/S0025-6196\(13\)00405-9/fulltext](https://www.mayoclinicproceedings.org/article/S0025-6196(13)00405-9/fulltext)

²⁴ <https://www.thenewatlantis.com/publications/why-data-is-never-raw>

²⁵ <https://www.technologyreview.com/s/612876/this-is-how-ai-bias-really-happensand-why-its-so-hard-to-fix/>

²⁶ <https://www.propublica.org/datastore/dataset/compas-recidivism-risk-score-data-and-analysis>

²⁷ <https://www.pbs.org/newshour/health/crispr-causes-significantly-greater-dna-damage-than-previously-thought-study-finds>

²⁸ <https://enterpriseproject.com/article/2019/1/ai-bias-9-questions-leaders-should-ask>

²⁹ <https://enterpriseproject.com/article/2019/1/ai-bias-9-questions-leaders-should-ask>

³⁰ <https://www.informationweek.com/big-data/ai-machine-learning/biased-ai-is-real-what-does-that-mean-for-women/a/d-id/1334223>

³¹ <https://enterpriseproject.com/article/2019/1/ai-bias-9-questions-leaders-should-ask>

³² <https://www.cio.com/article/3268965/new-ai-tools-make-bi-smarter-and-more-useful.html>

³³ https://www.dfs.ny.gov/industry_guidance/circular_letters/cl2019_01

³⁴ “DataOps is an automated, process-oriented methodology, used by analytic and data teams, to improve the quality and reduce the cycle time of [data analytics](#). While DataOps began as a set of best practices, it has now matured to become a new and independent approach to data analytics.”
<https://en.wikipedia.org/wiki/DataOps>

³⁵ <https://www.poolparty.biz/what-is-a-knowledge-graph>

³⁶ Robert Moore, “Using Graph Analytics for Predictive Modeling in Life Insurance,” Thrivent SAS,
<https://support.sas.com/resources/papers/proceedings17/1169-2017.pdf>

³⁷ <https://en.wikipedia.org/wiki/Chatbot>

³⁸ DevOps is a set of software development practices that combines software development (Dev) and information technology operations (Ops) to shorten the systems development life cycle while delivering features, fixes, and updates frequently in close alignment with business.

³⁹ <https://privacyblogfullservice.huntonwilliamsblogs.com/wp-content/uploads/sites/28/2019/04/AIEthicsGuidelinespdf1.pdf>

⁴⁰ “Human agency is the capacity for human beings to make choices and to impose those choices on the world. It is normally contrasted to natural forces, which are causes involving only unthinking deterministic processes. ... Human agency invests a moral component into a given situation.”
https://psychology.wikia.org/wiki/Human_agency

⁴¹ <https://ec.europa.eu/digital-single-market/en/news/policy-and-investment-recommendations-trustworthy-artificial-intelligence>

⁴² https://en.wikipedia.org/wiki/Black_swan_theory#Background:

About The Society of Actuaries

The Society of Actuaries (SOA), formed in 1949, is one of the largest actuarial professional organizations in the world dedicated to serving more than 32,000 actuarial members and the public in the United States, Canada and worldwide. In line with the SOA Vision Statement, actuaries act as business leaders who develop and use mathematical models to measure and manage risk in support of financial security for individuals, organizations and the public.

The SOA supports actuaries and advances knowledge through research and education. As part of its work, the SOA seeks to inform public policy development and public understanding through research. The SOA aspires to be a trusted source of objective, data-driven research and analysis with an actuarial perspective for its members, industry, policymakers and the public. This distinct perspective comes from the SOA as an association of actuaries, who have a rigorous formal education and direct experience as practitioners as they perform applied research. The SOA also welcomes the opportunity to partner with other organizations in our work where appropriate.

The SOA has a history of working with public policymakers and regulators in developing historical experience studies and projection techniques as well as individual reports on health care, retirement and other topics. The SOA's research is intended to aid the work of policymakers and regulators and follow certain core principles:

Objectivity: The SOA's research informs and provides analysis that can be relied upon by other individuals or organizations involved in public policy discussions. The SOA does not take advocacy positions or lobby specific policy proposals.

Quality: The SOA aspires to the highest ethical and quality standards in all of its research and analysis. Our research process is overseen by experienced actuaries and nonactuaries from a range of industry sectors and organizations. A rigorous peer-review process ensures the quality and integrity of our work.

Relevance: The SOA provides timely research on public policy issues. Our research advances actuarial knowledge while providing critical insights on key policy issues, and thereby provides value to stakeholders and decision makers.

Quantification: The SOA leverages the diverse skill sets of actuaries to provide research and findings that are driven by the best available data and methods. Actuaries use detailed modeling to analyze financial risk and provide distinct insight and quantification. Further, actuarial standards require transparency and the disclosure of the assumptions and analytic approach underlying the work.



Give us your feedback!
Take a short survey on this report.

[Click here](#)



Society of Actuaries
475 N. Martingale Road, Suite 600
Schaumburg, Illinois 60173
www.SOA.org