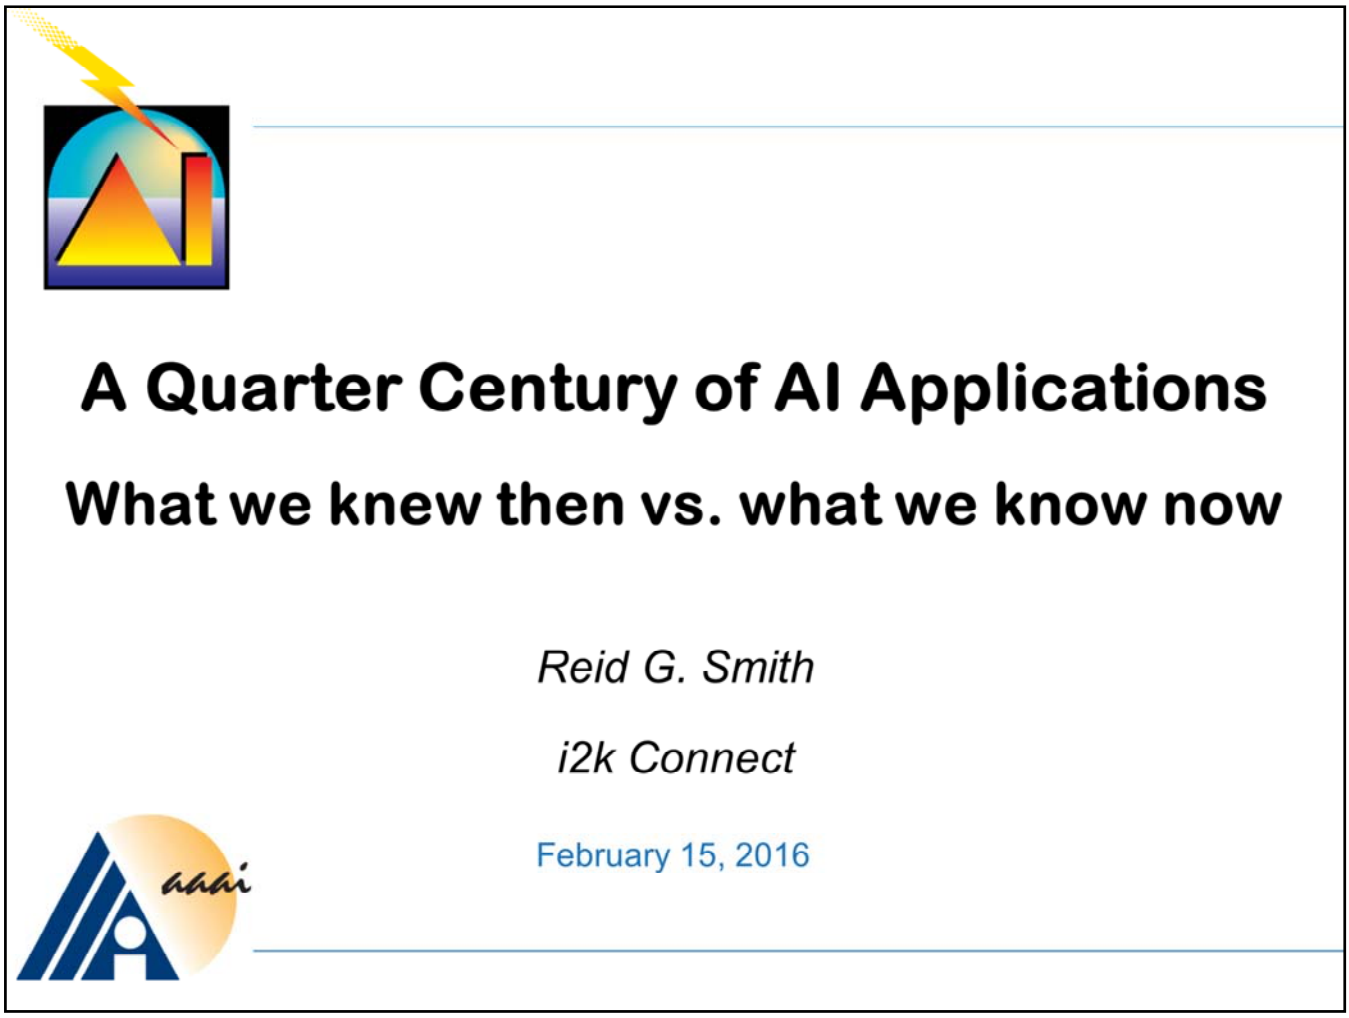

AI applications have been built, deployed and used for industrial and government purposes for many years. The experiences have been documented in IAAI conference proceedings since 1989. Over that period, the breadth of applications has expanded many times over. The diversity of technical approaches has also evolved — from rule-based expert systems to deep learning — with many modern systems employing a variety of techniques and subsystems. This presentation will focus on contrasting what (we thought) we knew about building, deploying and using AI applications in the early years with what (we think) we know now.

http://www.aaai.org/Conferences/AAAI/2016/aaai16speakers.php#Smith

First of all, let me say what a privilege it is to receive an award that honors Bob Engelmore's extraordinary service to AAAI and his contributions to applied AI.

Bob was a friend I first met as a Ph.D. student at Stanford, where he was Executive Director of the Heuristic Programming Project.

I hope these photos stir your memories as they do mine.

In the upper left, we see Bob with Penny Nii and Ed Feigenbaum. Some of you will remember that Bob and Ed were roommates at Carnegie Tech.

In the lower left are members of the 2000 IAAI program committee, in Austin: Neil Jacobstein, Ted Senator, Ramasamy Uthurusamy, Bob, Haym Hirsh and me.

In the lower right are Bob and Bruce Buchanan on one of their hiking adventures.

My only regret is that I could not find any photos of the tennis matches with Bob and Bruce ... Who could forget this line from Bob: "Too fast, Dave, take two."

Structure of the presentation

- AI Applications World – 1989: 1st IAAI conference
- AI Applications World – Today
- IAAI Data
  - Industry
  - Technology
- What (we thought) we knew then / what we learned
  - High-impact applications
- What (we think) we know now
  - What has changed
- What's next?

Today, I intend to tell you a story of the evolution of AI applications over a quarter century. The heroes of the story are our favorites – us – the AI community. It is a story about human-machine partnerships – the rise of the Intelligent Assistant as the dominant embodiment of AI in the workplace. It is also a story about the revolutions that have happened all around us – the explosion in computing power, the Web and open-source software – and how they have changed us.

Before we begin, the slides and the text will be available on my website: reidgsmith.com.

Time constraints have forced me to remove some of my favorite elements, but the follow-on AI Magazine article will incorporate them, as well as your comments, suggestions and corrections. So please don't be shy about sending them.

To cut to the chase,

What I want you to remember from the story is this: for AI to benefit humankind it has to be deployed; for successful deployment we have learned that good AI ideas have to be integrated into the human context of actual use and into the IT context of organizations.

AI Applications World – 1989: 1st IAAI conference

- **IAAI was the brainchild of Raj Reddy**

  ...and so was the AAAI, a decade earlier ...

- **1989: 1st IAAI conference – Stanford**

Like the AAAI itself, IAAI was the brainchild of Raj Reddy. It has been my privilege to be a Program Committee member for every conference.

Howie Shrobe summarized the context: "... the emergence of scientific achievements had triggered opportunities to tackle new problems ... The point of the conference was to exchange information about what really works and what the real problems are. The goal was to lead to better technology, to find and remedy current deficiencies, and to solve real problems."

IAAI-89 – Mostly Expert Systems

From Herb Schorr's introduction to the IAAI-89 book

- **Expert Systems**
  "Nearly all [applications] are expert systems because it is in this form that AI is most rapidly coming into widespread use."
- **Robotics**
  "no robot software system for complex tasks is commercially available ... robots seem to be stuck with their early applications and have made small commercial progress in the last few years.
- **Neural networks**
  "we know of no neural networks in practical day-to-day use ... while this technology appears to possess vast potential ... we leave it for this book's successor to cover such applications."
- **Natural language processing**
  "has been constrained historically by limitations of computational power, but the fantastic progression of computational cost/performance has eliminated this bottleneck. ... [But] today's applications ... are very limited and very few low-level natural language functions are being deployed."

The program chair, Herb Schorr, wrote in his introduction to the proceedings:

IAAI-89 – Mostly Expert Systems

From Herb Schorr's introduction to the IAAI-89 book

Expert Systems

"Nearly all [applications] are expert systems because it is in this form that AI is most rapidly coming into widespread use."

...

- **Neural networks**
  "we know of no neural networks in practical day-to-day use ... while this technology appears to possess vast potential ... we leave it for this book's successor to cover such applications."
- **Natural language processing**
  "has been constrained historically by limitations of computational power, but the fantastic progression of computational cost/performance has eliminated this bottleneck. ... [But] today's applications ... are very limited and very few low-level natural language functions are being deployed."

Expert Systems: "Nearly all [applications] are expert systems because it is in this form that AI is most rapidly coming into widespread use."

IAAI-89 – Mostly Expert Systems

From Herb Schorr's introduction to the IAAI-89 book

Robotics

"no robot software system for complex tasks is commercially available ... robots seem to be stuck with their early applications and have made small commercial progress in the last few years.

technology appears to possess vast potential ... We leave it for this book's successor to cover such applications."

- **Natural language processing**
  "has been constrained historically by limitations of computational power, but the fantastic progression of computational cost/performance has eliminated this bottleneck. ... [But] today's applications ... are very limited and very few low-level natural language functions are being deployed."

Robotics: "no robot software system for complex tasks is commercially available ... robots seem to be stuck with their early applications and have made small commercial progress in the last few years.

IAAI-89 – Mostly Expert Systems

From Herb Schorr's introduction to the IAAI-89 book

Neural networks

"we know of no neural networks in practical day-to-day use ... while this technology appears to possess vast potential ... we leave it for this book's successor to cover such applications."

- **Natural language processing**
  "has been constrained historically by limitations of computational power, but the fantastic progression of computational cost/performance has eliminated this bottleneck. ... [But] today's applications ... are very limited and very few low-level natural language functions are being deployed."

Neural networks: "we know of no neural networks in practical day-to-day use ... while this technology appears to possess vast potential ... we leave it for this book's successor to cover such applications."

IAAI-89 – Mostly Expert Systems

From Herb Schorr's introduction to the IAAI-89 book

Natural language processing

"has been constrained historically by limitations of computational power, but the fantastic progression of computational cost/performance has eliminated this bottleneck. ... [But] today's applications ... are very limited and very few low-level natural language functions are being deployed."

"has been constrained historically by limitations of computational power, but the fantastic progression of computational cost/performance has eliminated this bottleneck. ... [But] today's applications ... are very limited and very few low-level natural language functions are being deployed."

Natural language processing: "has been constrained historically by limitations of computational power, but the fantastic progression of computational cost/performance has eliminated this bottleneck. ... [But] today's applications ... are very limited and very few low-level natural language functions are being deployed."

Structure of the presentation

- AI Applications World – 1989: 1st IAAI conference
- **AI Applications World – Today**
- IAAI Data
  - Industry
  - Technology
- What (we thought) we knew then / what we learned
  - High-impact applications
- What (we think) we know now
  - What has changed
- What's next?

Let's fast forward to 2016.

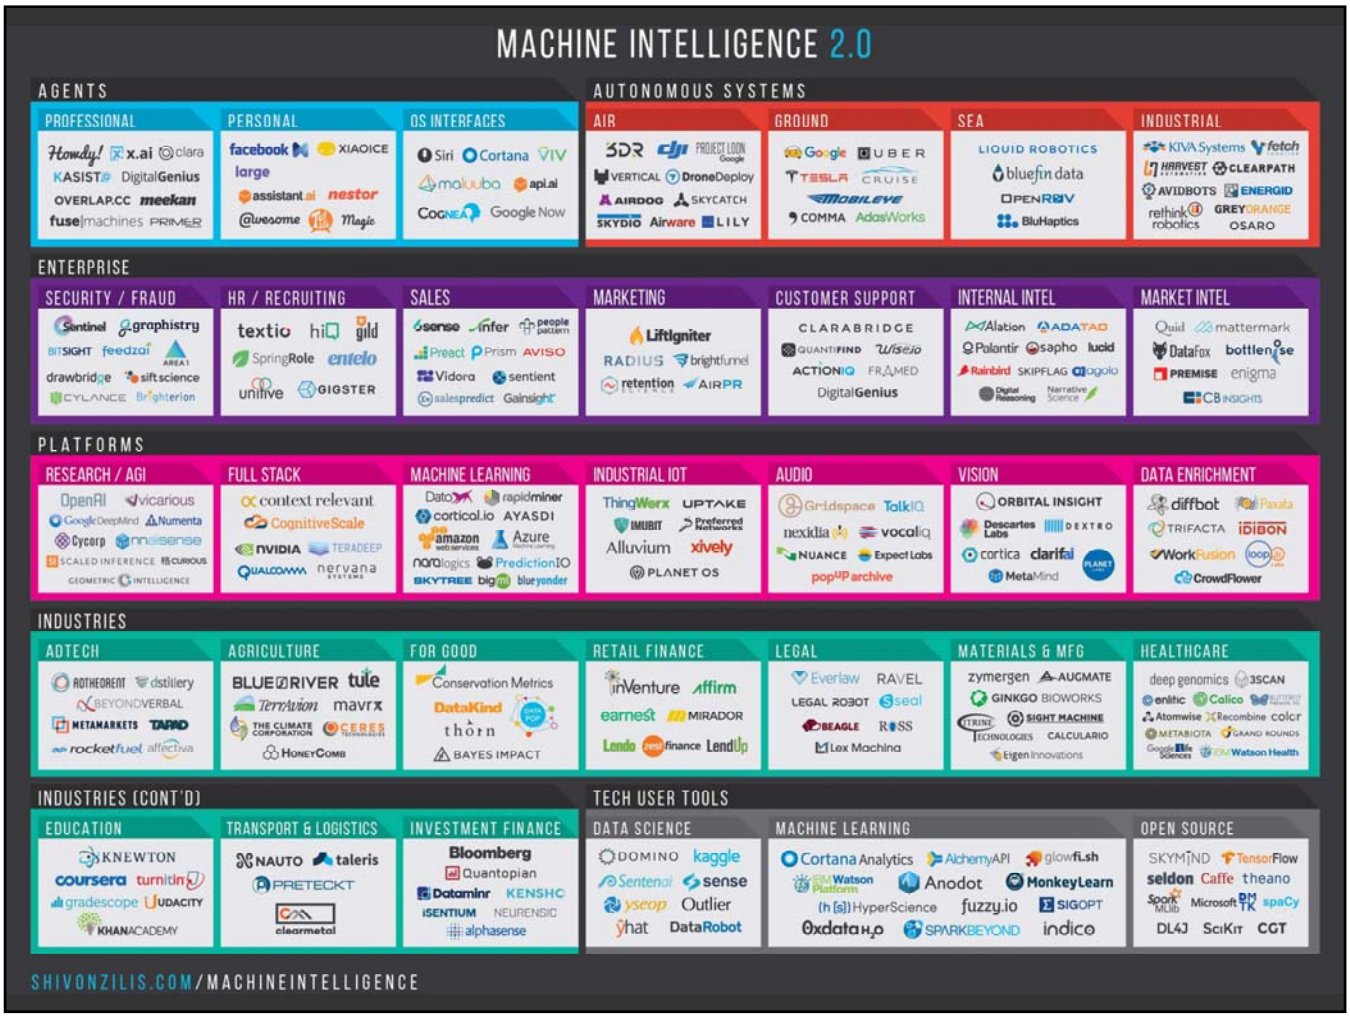

Today, there are so many companies working on AI that those of you beyond the first row will struggle to read the logos on this colorful chart from Shivon Zilis. Big and small players are putting a tremendous amount of money into AI technology.

She also makes an observation that is relevant to our discussion today: “startups [are] shifting away from building broad technology platforms to focusing on solving specific business problems.”

2016 – AI All Around Us

- Rule-Based Systems
- Credit Card Fraud Alert
- Insurance
- Scheduling
  - Maintenance, Crew, Gate, ...
- Video Games
- Search Engines
- Augmented/Virtual Reality
- Photo Face Recognition
- Handwriting Recognition
  - Mail Sorting, ATMs-Checks, ...
- Translation
- Deep Learning
- TurboTax
- Netflix Recommender
- FareCast, Google.com/flights, Kayak price predictor, ...
- Narrative Science GameChanger, ...
- Watson
- Dragon Speech Recognition
- Amazon Robotics / Kiva Systems
- Roomba
- Kinect
- Driver-Assist / Self-Driving Vehicles
- Siri, Cortana, ...

The world of AI apps is very different as well. In the early days, AI was viewed with suspicion in industry as the latest hype. Today, AI apps are all around us. Indeed, AI and machine learning are expected in almost every app.

This rather busy slide shows a few of the better known apps and technologies, many of which have been presented at AAAI or IAAI conferences over the years.

Of course not all are recognized as AI apps – the AI has disappeared into the fabric. Modern search engines are a good example of this phenomenon.

Note how many of the apps are Intelligent Assistants.

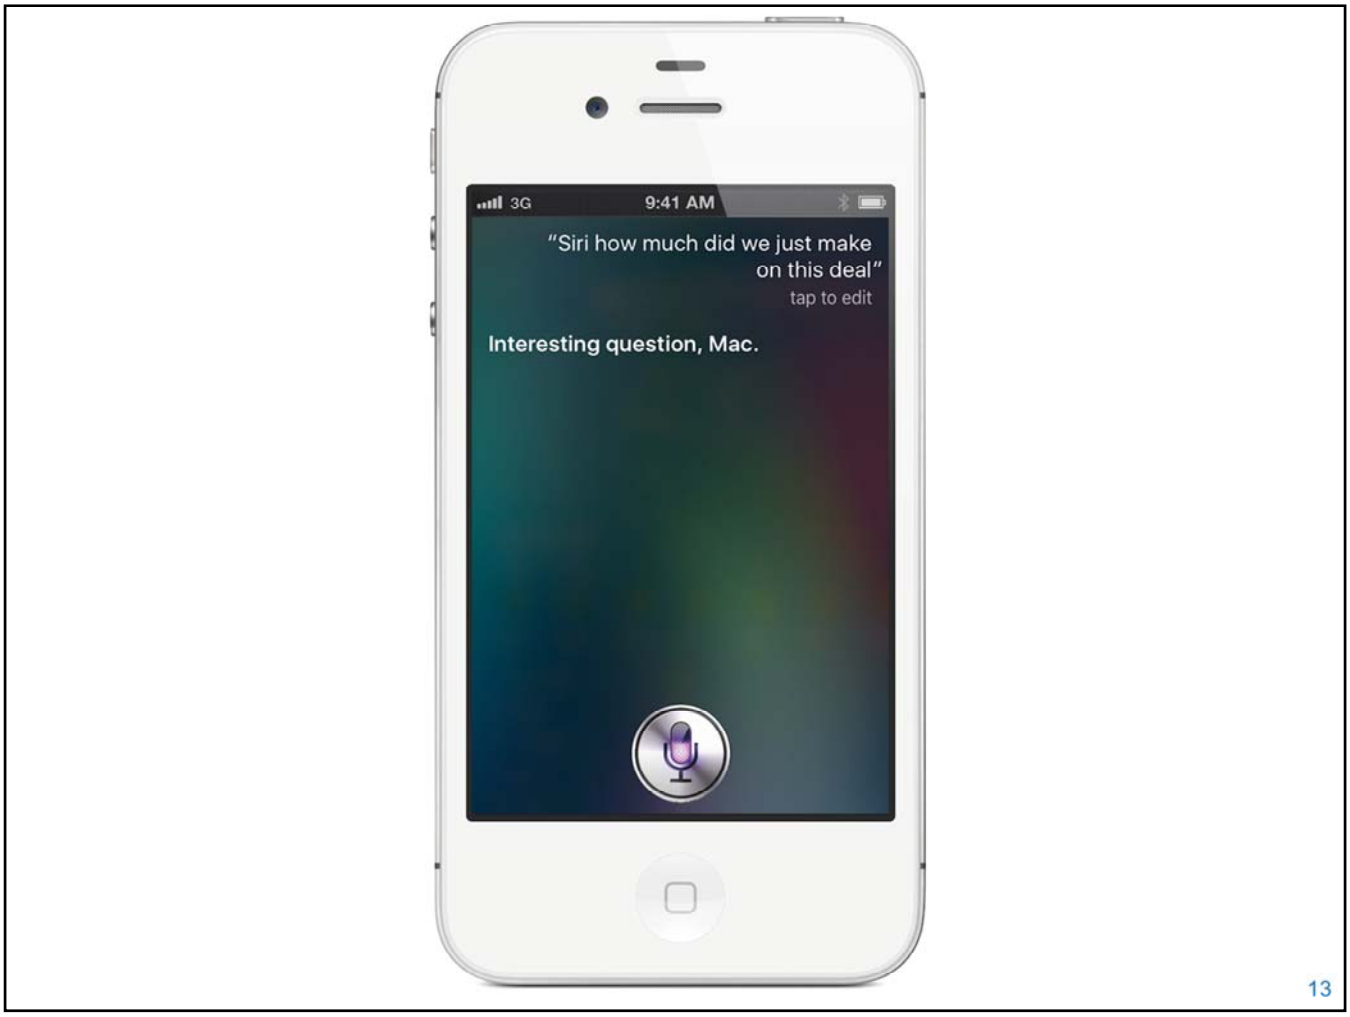

A little extra on Siri, from Adam Cheyer's invited talk in 2014.

Siri was introduced with the iPhone 4S in October of 2011. Headlines like this followed: "Siri provides 1 day return on investment as Apple nets $200 million day 1 of 4S launch."

I have a two-year old grandson. Because the iPhone button does not require much force to press and hold, he discovered Siri over a year ago. In the world of this child, he has always been able to talk with a machine and has always had an Intelligent Assistant.

Structure of the presentation

- AI Applications World – 1989: 1st IAAI conference
- AI Applications World – Today
- **IAAI Data**
  - Industry
  - Technology
- What (we thought) we knew then / what we learned
  - High-impact applications
- What (we think) we know now
  - What has changed
- What's next?

Next, let's look at distributions and trends from IAAI conferences.

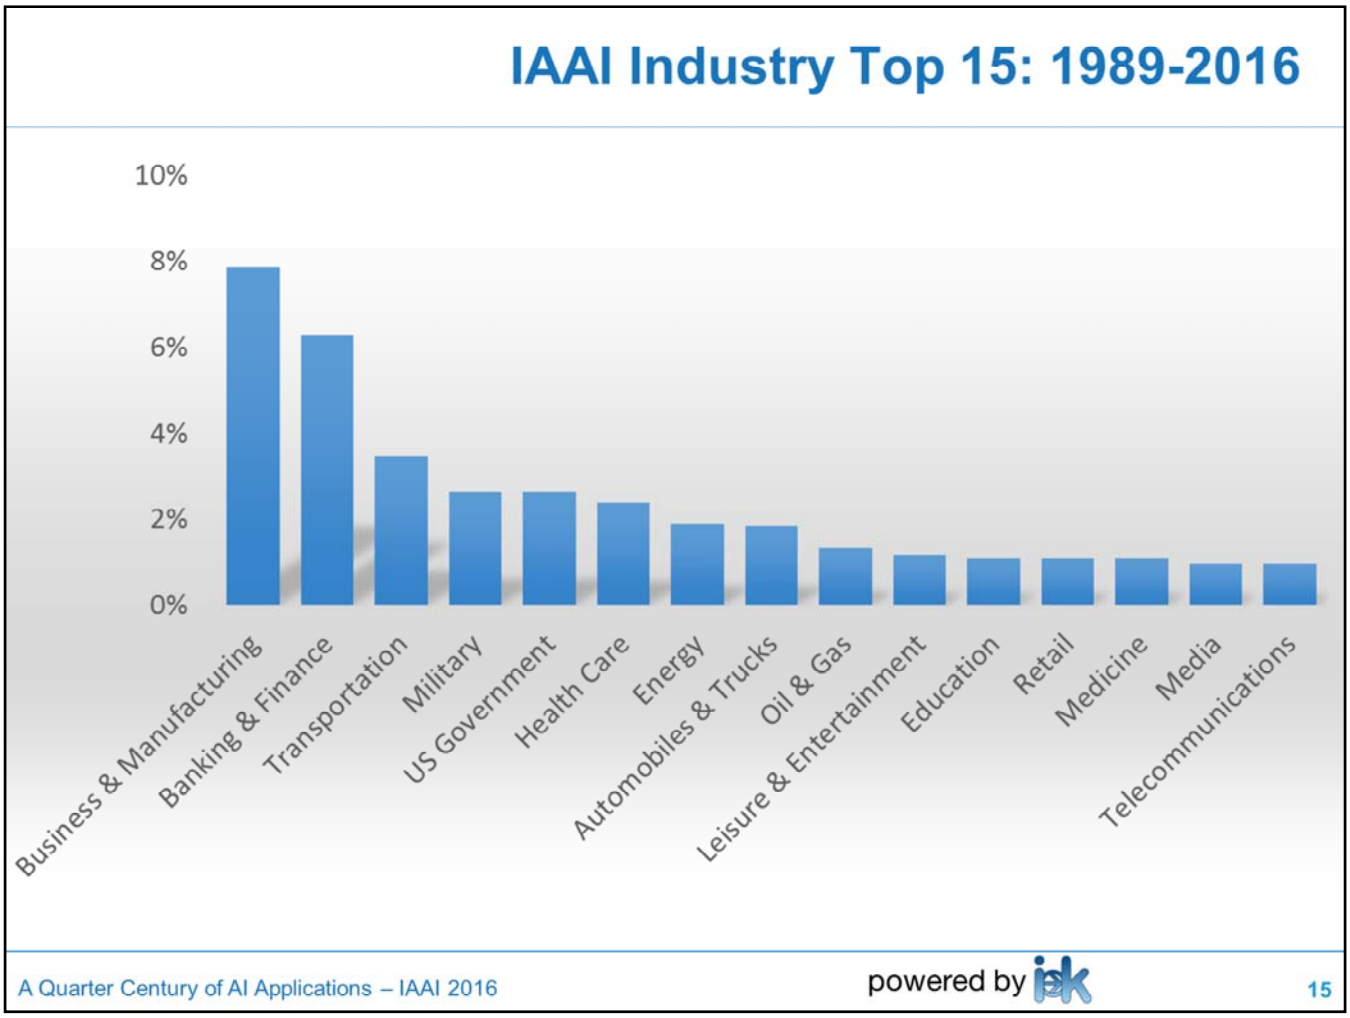

At the outset, IAAI included only applications that had been deployed; that is, for which there was experience based on use in the workplace, and for which payoff could be estimated.

In 1997, an emerging applications track was added to "bridge the gap" between research and applications.

This slide shows data from deployed applications only – 316 of them. We see the top 15 of a long tail of more than 100 domains that have been covered in IAAI. The chart excludes Information Technology apps; that is, AI applied to our own business, which would otherwise be number one.

A word about the analysis. It comes from i2k Connect. Our platform is an Intelligent Assistant that reads documents and automatically tags them with accurate and consistent metadata, guided and enriched by subject matter expertise. i2k Connect was founded by Bruce Buchanan and me. The other key members from the AI community are Eric Schoen and Joshua Eckroth.

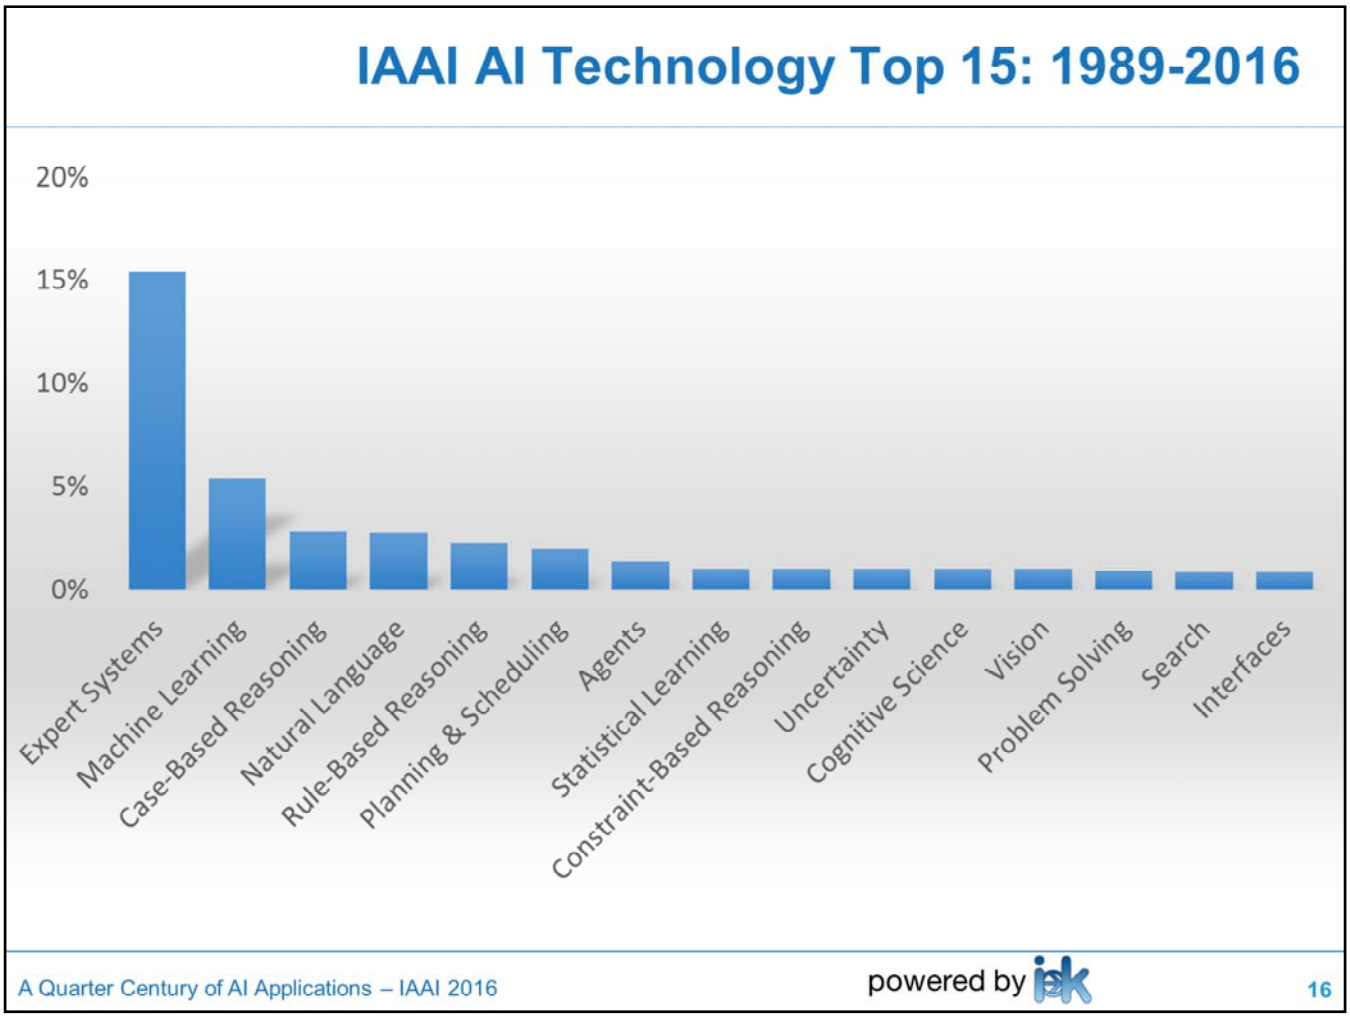

This slide shows the top 15 of a long tail of 85 AI technologies that have been incorporated into deployed apps.

A couple of things to note:

First, the Intelligent Assistant is the most common embodiment of the technologies in AI apps.

Second, it is interesting to note which technologies that haven't shown up as much at IAAI, like speech understanding and robots. They do show up, just not in the top 15.

Four technologies are applied in the 2016 Deployed Papers: Spatial Reasoning, Crowdsourcing, Machine Learning and Ontologies.

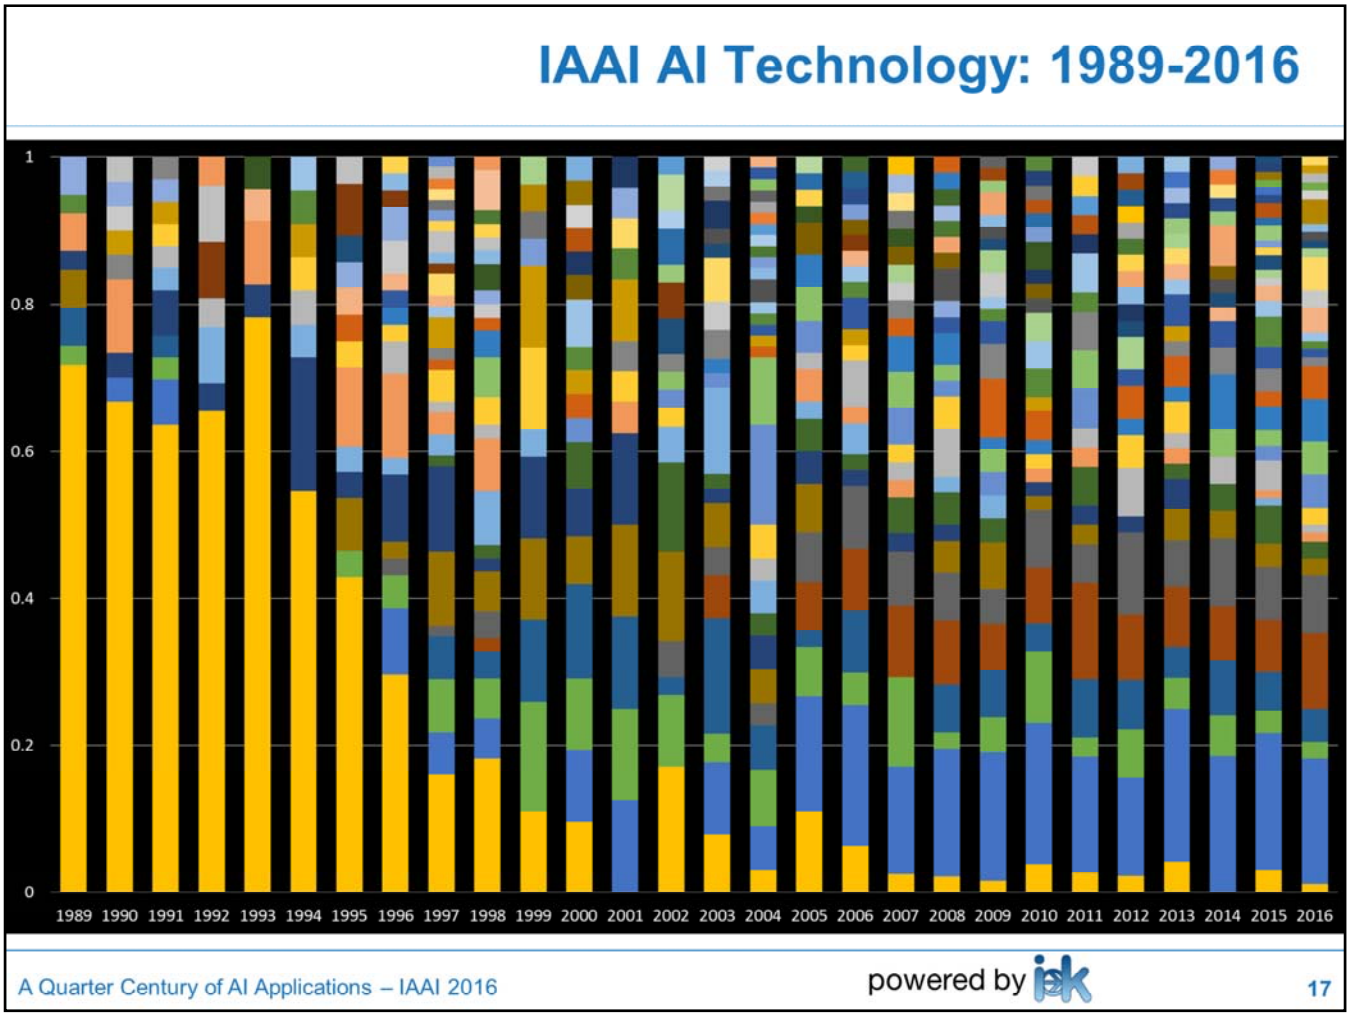

Here we see how the technology mix has evolved over the years. The slide includes data from all applications – deployed and emerging.

With the emerging apps included, the top five technologies are – start at the bottom: Expert Systems (yellow), Machine Learning (blue), Agents (green), Natural Language (darker blue) and Statistical Learning (rust).

You can see how dominant Expert Systems were in the early days. No surprise there. But note the absence of Machine Learning in 1989.

You may also be surprised by how diverse the mix has become over the history of the conference.

It is not the case that Expert Systems died. Rather, after a few years, they became more standard practice than innovation. They disappeared into the fabric, now applied everywhere, from the high-end emulation of rare human experts, to the embedding and application of rulebooks and procedure manuals.

Notes: There have been two papers on Deep Learning, one in 2015 and one in 2016, neither deployed.

Structure of the presentation

- AI Applications World – 1989: 1st IAAI conference
- AI Applications World – Today
- IAAI Data
  - Industry
  - Technology
- **What (we thought) we knew then / what we learned**
  - High-impact applications
- What (we think) we know now
  - What has changed
- What's next?

Now we have the data – a summary of our heroes' accomplishments. But every hero story must have challenges, obstacles and setbacks.

And so it is time to dive into what (we thought) we knew early on and what we have learned as the story of developing, implementing and upgrading apps has unfolded. Along the way, I will highlight a few high-impact applications – some that opened up a new area.

Because I have selected high-impact applications and it takes time to establish whether an application has had high impact, the examples may look a bit dated. Note, however, that in several cases, a recent update has been presented at IAAI.

Another hint: Intelligent Assistants play a big part what is to come.

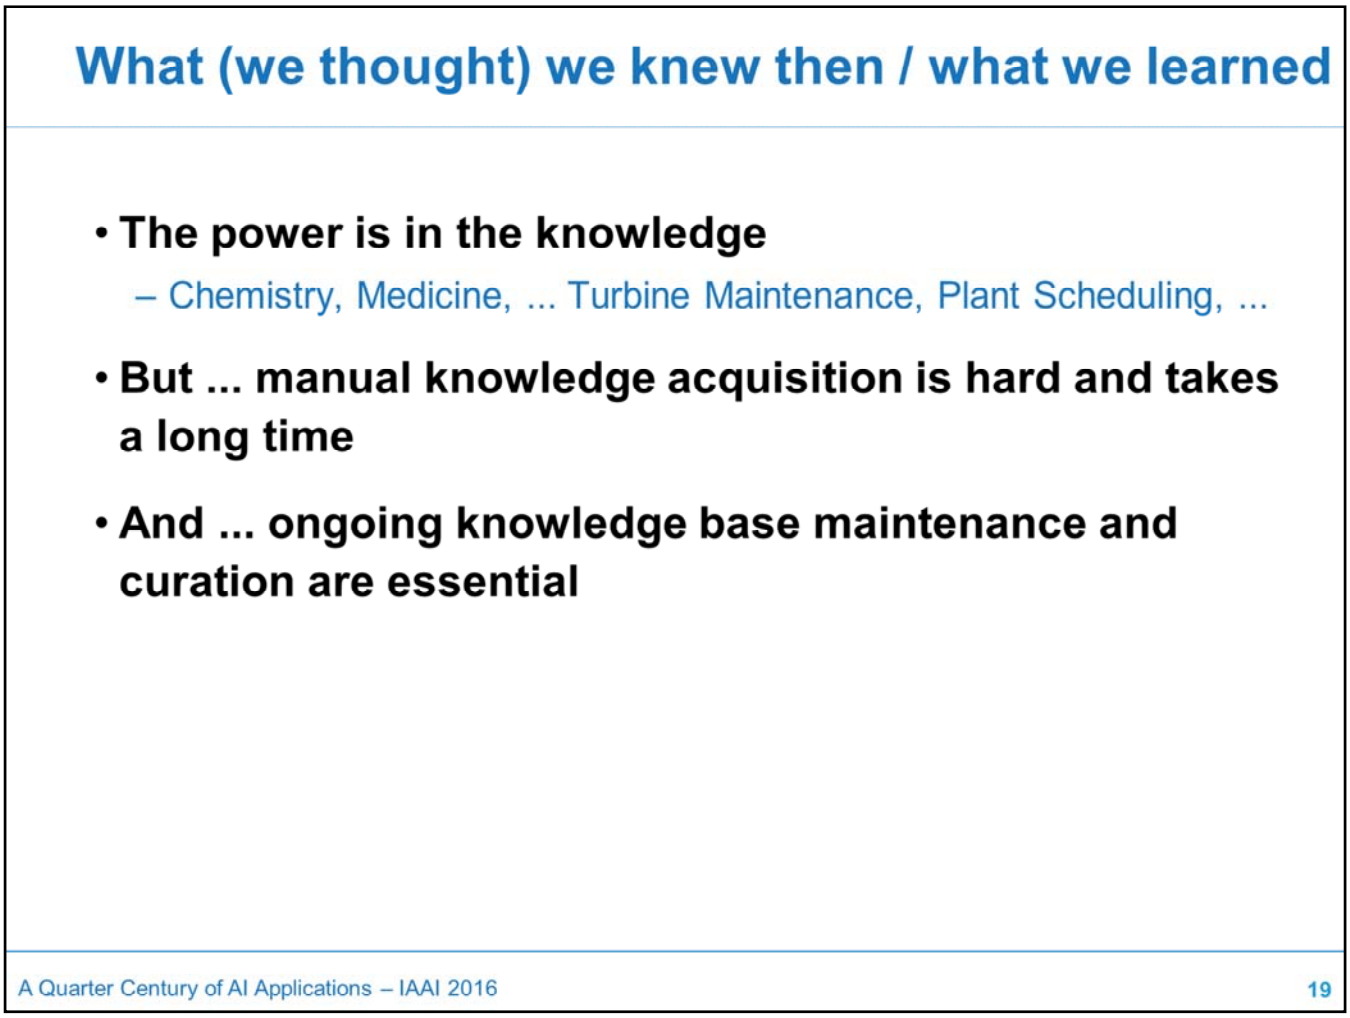

By 1989, thanks to the pioneering efforts of Ed Feigenbaum, Bruce Buchanan and many others, we knew that “the power is in the knowledge”; that is, domain-specific knowledge (chemistry, medicine, ...) and task-specific knowledge (turbine maintenance, plant scheduling). We understood that specific knowledge is more important for high performance than general problem-solving approaches.

So ... job done, right?! Just build the knowledge bases and away we go.

But ... And there is always a “but” ... Manual knowledge acquisition is hard and takes a long time. Indeed, it is the number one obstacle in the history of AI apps – the well known knowledge-acquisition bottleneck.

And ... worse yet, ongoing knowledge base maintenance and curation are essential. Why? Knowledge is perishable -- everything changes over the lifetime of an app: the domain evolves, new use cases arise, new experts arrive with different knowledge, the technology advances, and so on.

So, with manual knowledge acquisition, you need an army of people to keep the knowledge base up to date; and therefore, you need a lot of revenue and users to support that effort.

As a result, many systems fall by the wayside, even if they were excellent at one time. Take the case of INTERNIST and CADUCEUS, the work of Harry Pople, Randy Miller and Jack Myers. At one time, perhaps the most knowledge-intensive systems in existence: 500,000 pieces of internal medicine knowledge. It could diagnose 1,000 diseases and performed beyond what people could do at the time. But efforts to commercialize it failed, and without a revenue stream, it could not be maintained and is no longer in use.

1985: Process Diagnostics System

- **Started as Expert System Shell**
- **Evolved through many changes: technology, regulatory environment, company, ...**
- **30 year anniversary paper in 2015**

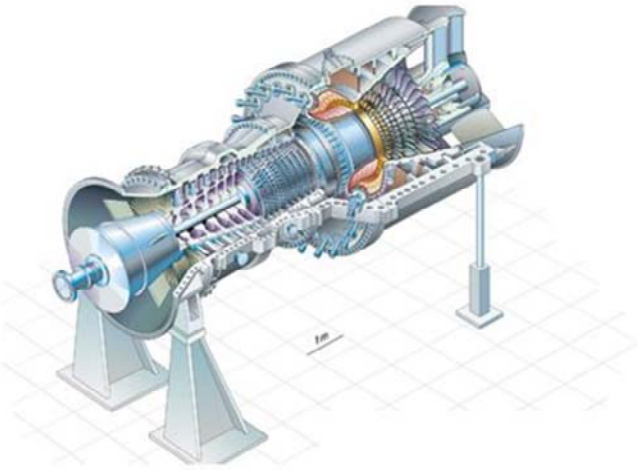

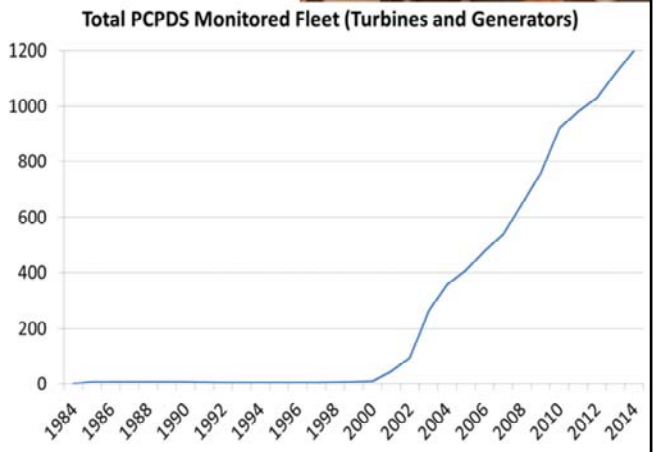

With that in mind, let's take note of PDS, the Process Diagnosis System. It has been in active use and continuous development since 1985. **1985!** It started with a presentation by Mark Fox at GE. Over the 30 year period, GE sold the business to Siemens, where it is now at the heart of their Power Diagnostics® Center that performs centralized rule based monitoring of over 1,200 gas turbines, steam turbines, and generators.

Ed Thompson and Ben Bassford celebrated the 30th anniversary with us when they presented an update at the 2015 IAAI in Austin.

Knowing the challenges that must be overcome for ***any*** application to see the light of day in industry, I have utmost respect for the PDS team.

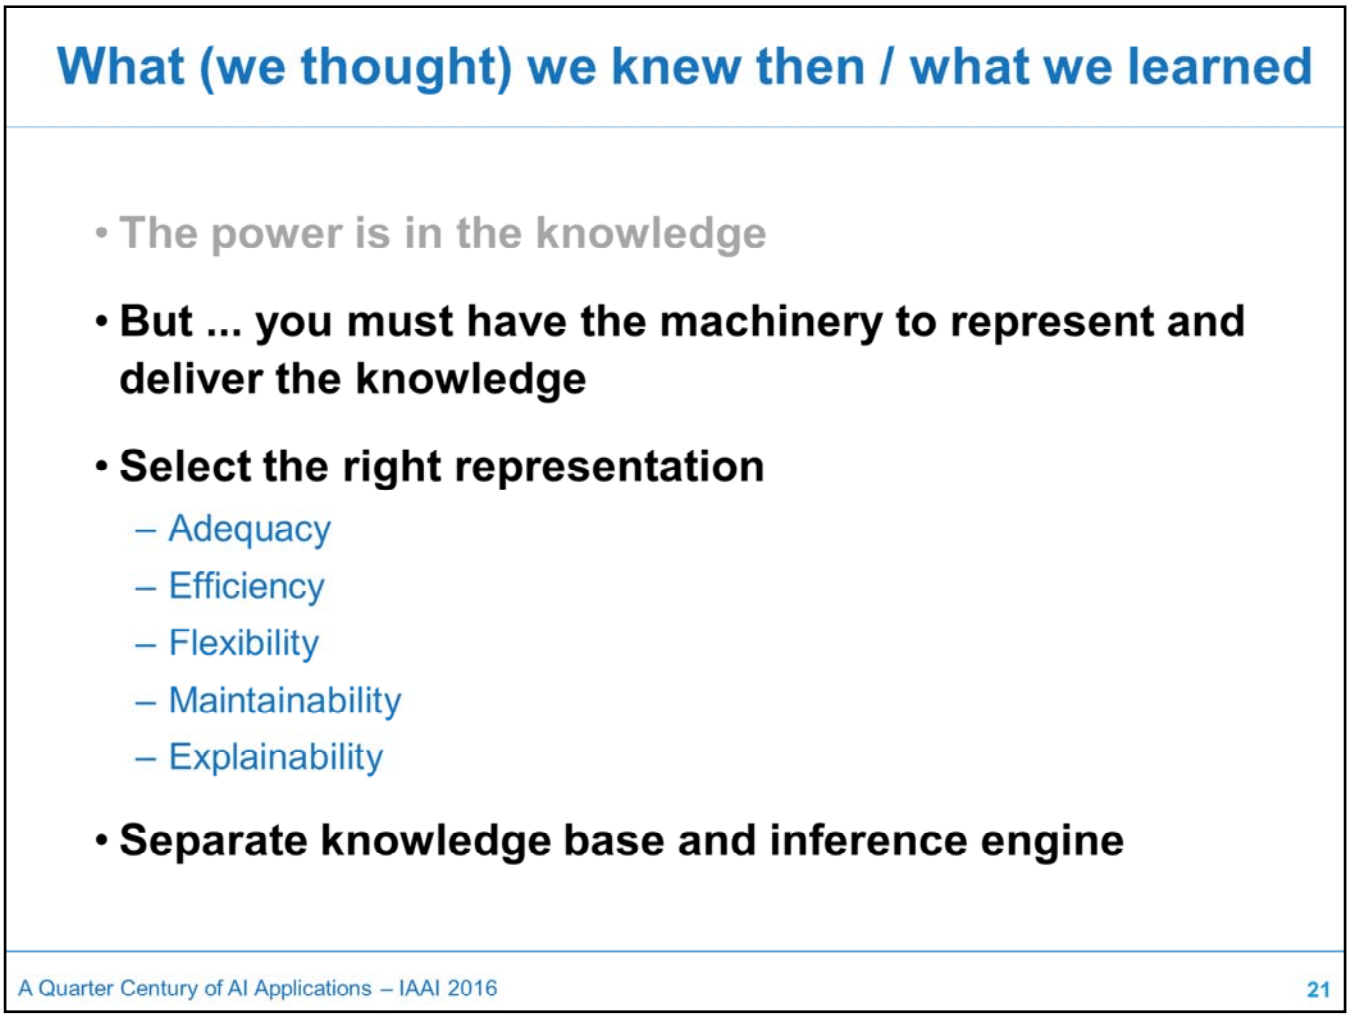

What did our AI apps heroes learn next? Another But! You must have the machinery to represent and deliver the knowledge.

The pioneers taught us that selecting the right representation has a big impact – *matching the representation to the domain and task characteristics*. Here are some of the reasons.

John McCarthy and his colleagues taught us about epistemological adequacy – you can't reason about what you can't represent.

Every programmer – numeric or symbolic – knows that selecting the right representation can have an enormous impact on efficiency.

The other day, Ted Senator reminded me of the impact of representation on the next three items.

Flexibility: If you make a good representation choice, you may be able to use your app to address problems you didn't anticipate at the outset.

Maintainability: It is very useful for subject matter experts to be able to maintain the knowledge base without the need to reach out to AI geeks. Another McCarthy lesson – declarative representations are more learnable and maintainable than procedural ones.

Explainability: And finally, remember the Intelligent Assistant theme. An assistant is expected to be able to explain the reasoning – to a variety of people – end users, subject matter experts, system developers and the management. Why the management? Well, the knowledge base represents the decision criteria of the company. It is called a "knowledge asset" by knowledge management

practitioners. Therefore management must own it and be able to defend it.

Lately, there has been some discussion in the press about “algorithmic accountability” and several companies are pursuing explanation as a differentiator (e.g., Watson Paths and the Narrative Science extension for the Qlik visual analytics tool).

The pioneers also taught us that it is a good idea to separate the knowledge base and the inference engine. Why? It is easier to change, update, debug, and explain.

These two points – representation and delivery – sparked a tremendous amount of work on Expert System shells (like M1, S1, ART and CLIPS), knowledge representation languages (like KL-ONE and OWL), and ontology editors (like Protégé).

1995: FAIS & 1998: ADS

- **FAIS: FinCEN Artificial Intelligence System**
  - Money Laundering
  - Link Diagrams
- **ADS: NASD Regulation Advanced-Detection System**
  - Securities Fraud
  - Temporal Sequences
- **Different domain concepts dictate different representations**

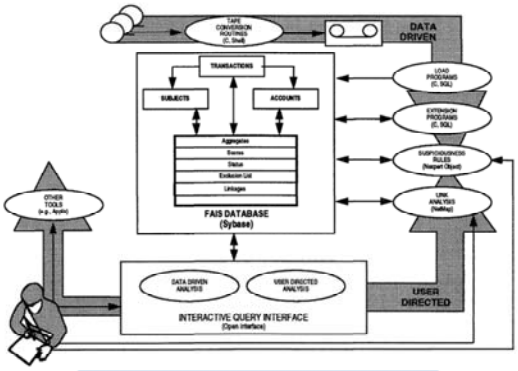

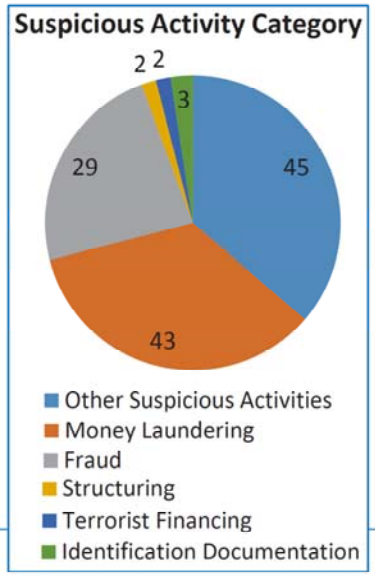

The two systems shown here illustrate these points.

Two more Intelligent Assistants! The work of Ted Senator and his colleagues.

FAIS links and evaluates reports of large cash transactions to identify potential money laundering.

To give you an idea of the money involved, the diagram on the right is from fincen.gov, October 2015. The total reported transaction amount is approximately $28B.

The key idea for FAIS is connecting the dots ... so a representation that supports the link diagrams now commonplace in social network analysis is a good choice.

ADS monitors trades and quotations in the Nasdaq Stock Market, to identify suspicious patterns and practices. In this app, temporal sequences are key – not so much links as in FAIS – so a representation that supports them is a good choice.

And speaking to flexibility, the right representation choice in FAIS enabled a reporting app based on the original detection system. This was an unanticipated bonus.

What (we thought) we knew then / what we learned

- The power is in the knowledge
- **And ... successful AI applications incorporate a variety of knowledge and techniques**
  - Rules
  - Statistics
  - Signal Processing

    ...

Our heroes also learned that successful AI applications incorporate a wide range of techniques and strategies, embodying rules, objects, ontologies, statistics and signal processing to name a few.

Self-driving cars are an obvious example.

Their capabilities include – modeling, simulation, sensing, motion planning, object recognition, obstacle avoidance, machine learning, flexibility, error recovery, and so on.

What (we thought) we knew then / what we learned

- The power is in the knowledge
- **But ... you must have the machinery to integrate into the workflows that people actually use**
- **So ... your machinery must play nicely with other tools**
- **And ... it must play nicely with people**

The Intelligent Assistant ... again

Oh oh ... Another “but”. Our heroes learned that success depends on integrating into existing workflows – the human context of actual use. It is rare to completely replace an existing workflow.

This means that the AI components must play nicely with the other tools that people use. The AI is a piece of the puzzle, and sometimes not a very big piece of the total effort.

And at the outset, the business and technical infrastructure barriers were large. On top of that, AI exploitation required boutique hardware and software, as well as scarce developers.

And we learned how important it is to play nicely with people – easy to use, able to explain the reasoning, and actually solving a problem they have – regardless of whatever our favorite AI tool might be.

Put another way, the human interface is the “license to operate”. Unless you get that part right, people may not ever see the AI power under the hood. They will have already walked away.

Working in the real world is messy. The data are messy, the IT systems you have to integrate with are messy, and the people are messy.

And once again, but not for the last time, most AI systems play the role of the Intelligent Assistant in a human-machine partnership.

In the real world of applications, there is no dichotomy between Artificial Intelligence and Intelligence Augmentation or Amplification. They are two ends of a spectrum that meet in most applications. The successful systems enable people do what people do best and use computers to do what computers

do best.

An example from the early days was “the miracle of the Dendral algorithm”. It was devised by Nobel Prize winner Joshua Lederberg for exhaustive and non-redundant enumeration of the large space of molecular structures that obey structural constraints. It was the legal move generator for the hypothesis space of chemical structures and the heart of the Plan-Generate-Test paradigm.

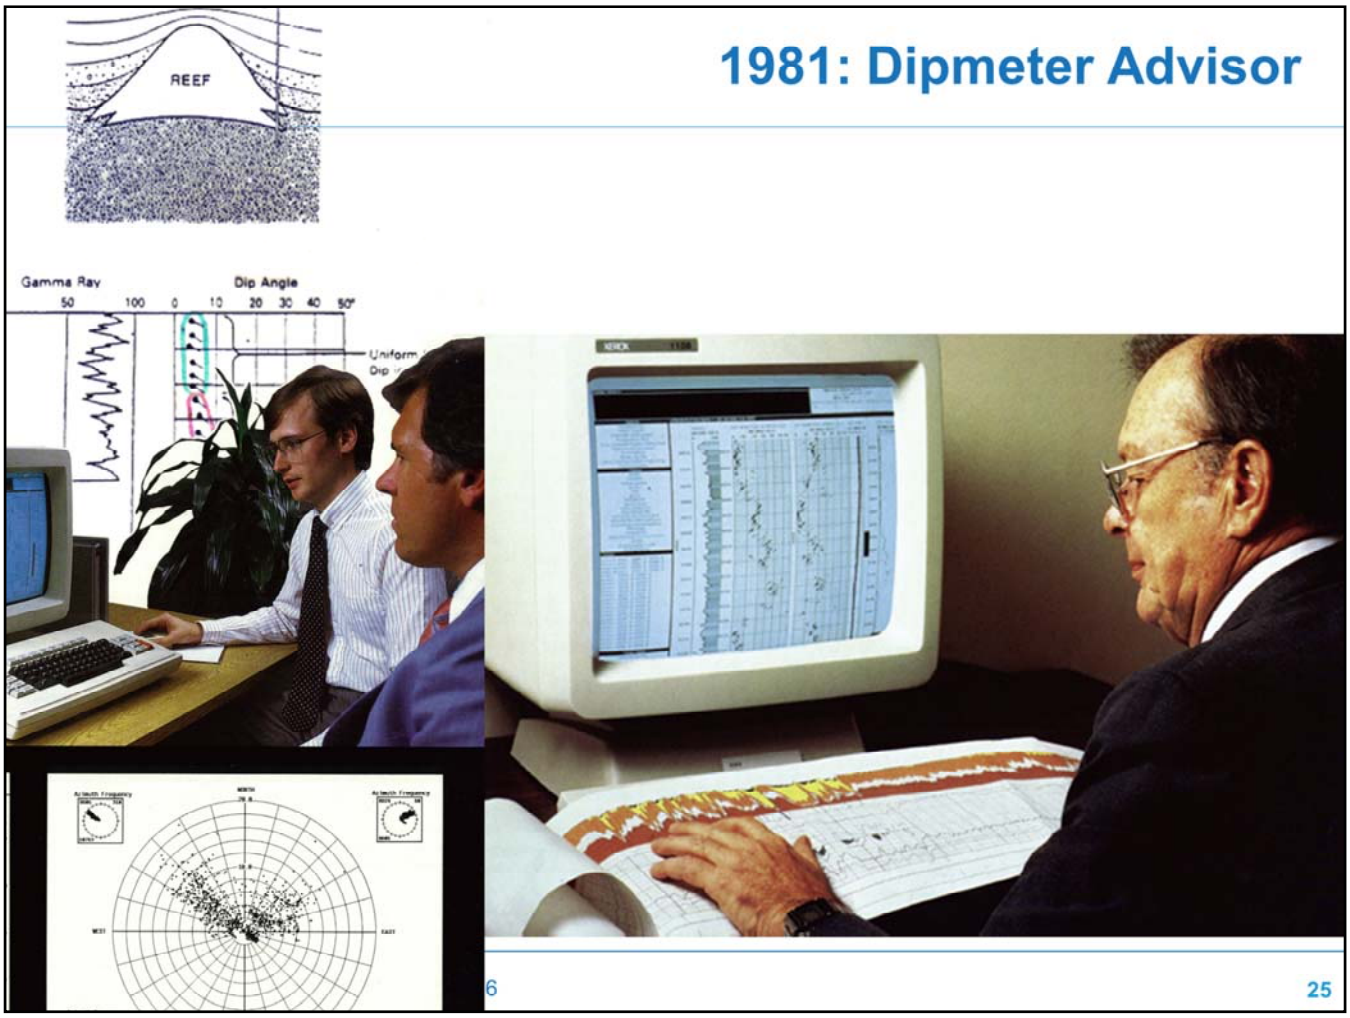

Work on the Dipmeter Advisor started at Schlumberger in the early 80's, based on the knowledge of the legendary oil finder, Al Gilreath (shown in the lower right). Also shown here is Dave Hammock (in the white shirt), who made important contributions to building, applying and selling the technology to clients.

The Dipmeter Advisor demonstrated the challenges of infrastructure (e.g., getting the data from the field systems was a bigger problem than I anticipated); and technology transfer (boutique hardware – D-Machines – and software – Interlisp-D – were major stumbling blocks – but without them we would have had no system at all).

As in many systems, the amount of effort that had to be devoted to the non-AI components was dominant. For the Dipmeter Advisor, the user interface accounted for almost half the code. The rule engine and knowledge base accounted for 30%.

(The whole story: Inference Engine: 8%; Knowledge Base: 22%; Feature Detection: 13%; User Interface: 42%; Support Environment: 15%).

Of course lines of code don't necessarily tell the whole story, but the numbers are consistent with the development effort expended. A lot of my time went into the interactive graphics system – and for some clients, that was the most important element – not the AI.

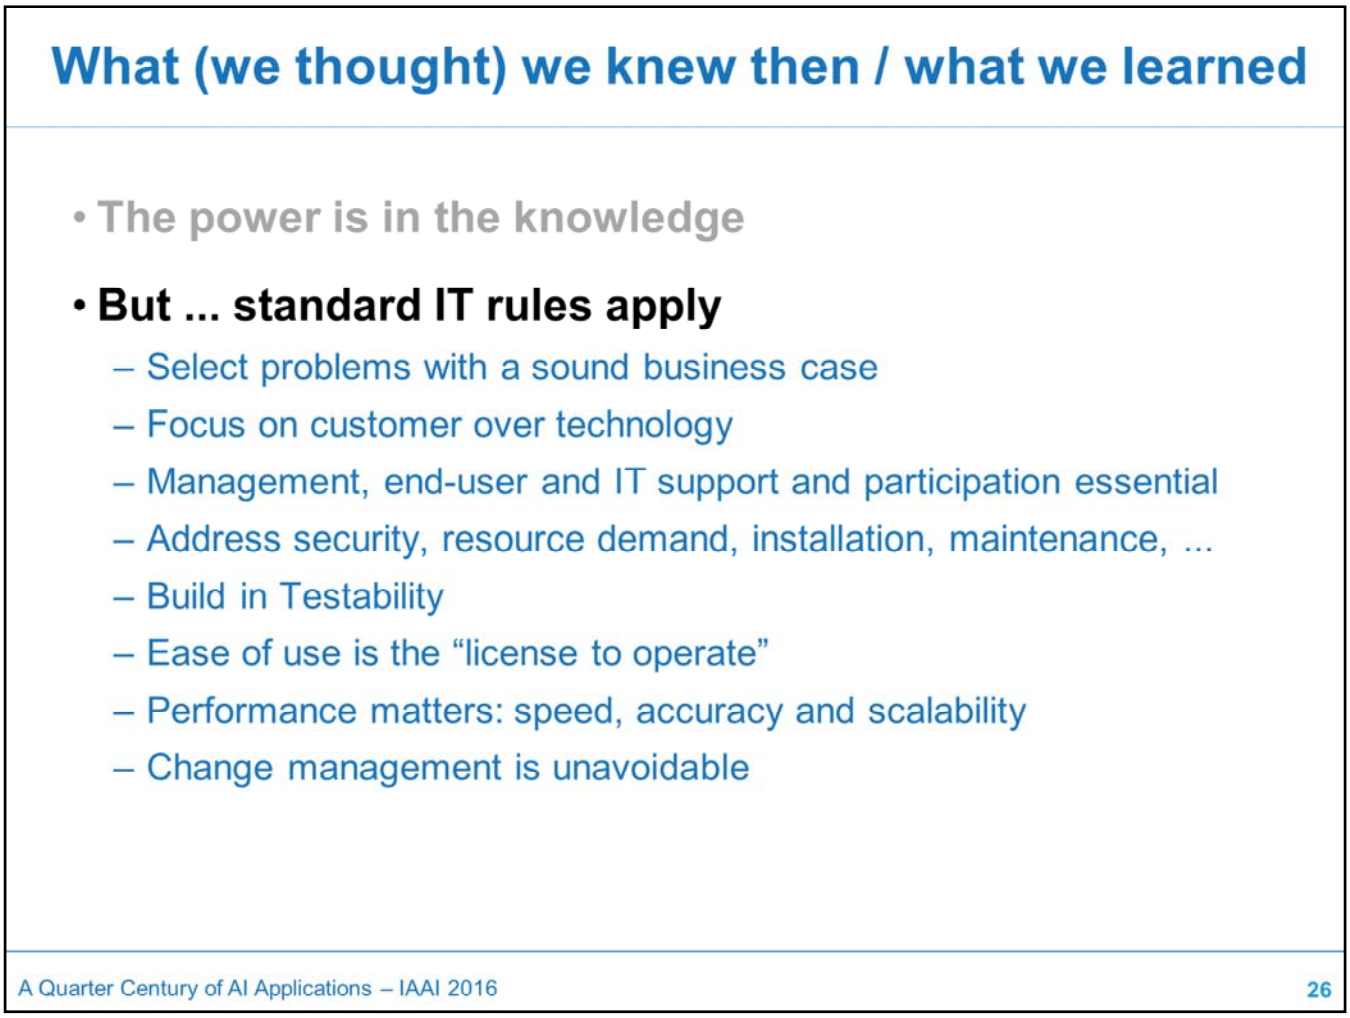

As the systems built by our heroes started to play in the center ring, so to speak, risk mitigation, project management and budgetary control became more important. The systems were no longer in a "research" or "proof of concept" phase. In other words, standard IT rules apply.

Many AI practitioners have made these points in the context of AI applications in particular. But "AI" is a variable in the phrase "AI applications". It can be replaced by any information technology and the rules still hold.

Select problems with a sound business case – don't waste your time and money.

Focus on customer over technology – our heroes made a major mind shift from the early days in this regard.

Management, end-user and IT support and participation are essential – you won't get budget approval without management support and you won't succeed without all of these elements.

Address security, resource demand, installation, maintenance, ... Security and privacy have become increasingly crucial over time.

Build in Testability. This is especially important, given the amount of KB maintenance that will be done over time.

And as I noted previously, ease of use is the "license to operate". But it isn't enough by itself.

Performance matters: speed, accuracy and scalability – and research speed and accuracy standards

are not always as rigorous as industry or consumer expectations. To get good numbers, you have to deal with all of those messy special cases.

Change management is unavoidable – The amount required is inversely proportional to the power of the new technology and the amount of change required to adopt it. As Mehmet Goker puts it: “Applications with a small and flexible core that solve a real world problem have the biggest impact and are the easiest to put into the workplace.”

AI Apps Central to Business

Ford Motor Company Direct Labor Management System

– Automatic generation of work instructions with associated times

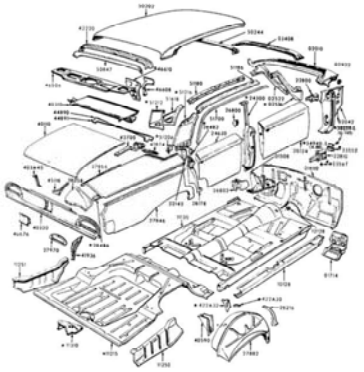

Engineering Works Scheduling for Hong Kong's Rail Network

– 5M people per day

Here are two Intelligent Assistants that played in center ring from the get-go. Both have been presented twice at IAAI – covering initial deployment and a 10 year update.

The first app is integrated into Ford's Manufacturing Process Planning System. The app has a natural language component, one of the few from the early days. The big payoff is automatic generation of work instructions for vehicle assembly, with associated times.

The second app enables the Hong Kong rail network to move 5M passengers a day through the city's rapid transit subway, airport express and commuter rail lines. The system must interact well with staff members and be able to explain the schedules it creates. As a result, over time the team veered away from Genetic Algorithms towards Heuristic Search.

Thanks to Nestor Rychtyckyj and Andy Chun for sharing their experiences. These two are among the IAAI deployed apps champions.

Structure of the presentation

- AI Applications World – 1989: 1st IAAI conference
- AI Applications World – Today
- IAAI Data
  - Industry
  - Technology
- What (we thought) we knew then / what we learned
  - High-impact applications
- **What (we think) we know now**
  - What has changed
- What's next?

It is time to summarize what (we think) we know now about building AI apps. But to put the changes in perspective, it is mandatory to start by addressing the revolutions we have seen over a quarter century.

What has changed?

- **10,000 times more computing power**
  - Commodity hardware
- **The Web**
- **Open-source software libraries**
- **The Cloud**
- **Machine learning practical for many problems**
- **Robotics and Natural Language Processing arrived**
- **Successful high-profile AI systems all around us**

I am not telling you anything you don't already know, but I do think it is important to remind ourselves that the key influencers have come from outside our AI world.

First, the ongoing computing revolution: From 1989, this has improved performance by a factor of 10,000! (1,000 in single thread performance x 10 cores per processor)

And another factor of 128 or so to cover the increase in memory size (128KB to 16GB).

And another factor of 50,000 or so to cover the increase in disk size. (40MB to 2TB).

And then there is the performance of today's massively parallel GPU's: 6TF and beyond in a card!

Add to this the fact that many apps no longer need to run on boutique hardware.

And the Web revolution, starting from almost no content in February of 1993, when the Mosaic browser was released at NCSA, to 5B indexed pages as of today. This has made crowdsourced data of all sorts available for training and mining (structured, unstructured, labeled, and unlabeled).

And then another human revolution: open-source libraries, available on the Internet for almost any computation you can imagine.

Drawing from Francis Bacon and Bob Metcalfe in my knowledge management days, I used the equation: power = knowledge^{shared} – the power of knowledge goes up exponentially with how broadly it is shared. I believe it today even more than I did then.

When I look back, I am amazed that the pioneers were able to field apps at all in the early days, given

how little computing power and storage were available to them – and no Web!

And we are today witnessing the leading edge of another revolution – the Cloud! You can now bring to bear as much computation, storage, and network infrastructure as you need. It is no longer an IT infrastructure obstacle. It is (only!) a monetary issue.

I spent enough time as a research director to remember one recurring phrase from time spent with scientists in the lab: “but wait, there’s more” ... And once again there is more.

Building on the revolutions in computing and the Web, the research community – AI and beyond – developed new machine learning algorithms that are now practical for many problems. It changes everything when you have billions of examples and the power to process them.

When I say “the research community – AI and beyond”, I include a large community of engineers. Let me single out Bernie Widrow at Stanford as an exemplar. He is a EE professor who has done a great deal of machine learning work. He also holds a special place in my heart because he was instrumental in enabling two Stanford EE students – Tom Mitchell and me – to do our research at the Heuristic Programming Project in the CS department.

Add to this list the tremendous progress made in Robotics and Natural Language; two areas that were not yet mature enough for many AI applications in 1989.

Finally, the success of high-profile AI systems such as Watson and Siri means that almost everybody knows that AI can work in the real world. This has lowered the barriers to entry for AI applications in industry.

What (we think) we know now

- It is very hard to build and maintain large knowledge bases by hand
- But ... data-driven, statistical methods are "unreasonably effective" in several domains
- Machine learning is a way around the KA bottleneck
- But ... success often depends on human insight folded into the methods, like the choice of features

Now that our heroes have accumulated a quarter century of AI apps experience, overcome many obstacles and lived through several revolutions, how has the generally accepted knowledge changed?

First of all, one thing has not changed. It is still very hard to build and maintain large knowledge bases by hand. The manual knowledge-acquisition bottleneck is firmly in place.

But now armed with billions of crowdsourced examples from the Web, we have learned that data-driven, statistical methods are "unreasonably effective" in several domains. The statistics bring the ability to deal with noise and to cover problems where humans either have difficulty explaining how they do it, or where they don't do it very well in the first place.

The bottom line is that machine learning is a way around the Knowledge Acquisition bottleneck in a surprisingly broad number of domains.

But ... Howie Shrobe made an observation that rings true for me. "... when you look closer at successful statistical approaches, a lot of the success is in the choice of features to attend to or other similar ways of conveying human insight to the technique ..."

Remember also that we don't have billions of examples for every problem.

What's Next?

- **Trends continue in computers, data, ... sensing**
- **Apps increasingly data-driven, guided by human knowledge**
- **Intelligent Assistants improve quality of life**
  - Unstructured data no longer lost in the cloud
  - The end of Clutter overload

We are almost at the end of the story that has been written to date.

As promised, here are a few points on what I think is likely to happen next – not in the long term – I'm no visionary – but over the next few years.

First, the revolutions will continue all around us, in computers and data, as well as sensing.

So it follows that apps will continue to get more powerful and knowledgeable. They will cover a broader array of domains and tasks.

It also follows that apps will be increasingly data-driven, guided by human knowledge. And they will have a lot more data available, as the Internet of Things takes off.

Finally, Intelligent Assistants will be even more proficient at improving quality of life. The partnership between human and machine is going to be stronger and closer.

How will they improve quality of life? Jobs are more satisfying when we are able to focus on the "real" work. Intelligent Assistants will deal with the clutter of low-level tasks, or tasks that require extended concentration, consistency, scale, and so on.

In my little world of i2k Connect, there are big opportunities with unstructured data. It will no longer be lost in the cloud – whether the corporate cloud or the Internet cloud. We will have the tools to find it and unlock its connections. We will also have the tools to extract the essential information from the cluttered real-time streams that overwhelm us today.

What the AI Community Can Do

- Continue to expand the range of the possible
- Take Intelligent Assistants and other AI apps to the next S-curve

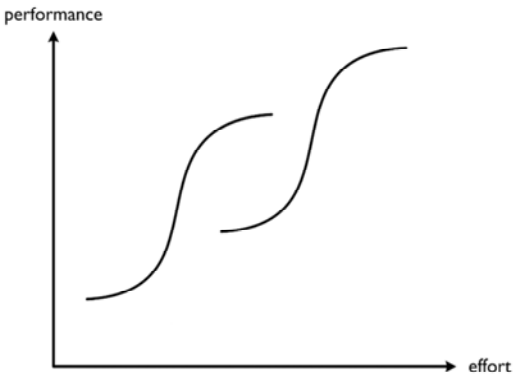

For AI to benefit humankind it has to be deployed

For successful deployment, good AI ideas have to be integrated into the human context of actual use and into the IT context of organizations

Finally, a few thoughts on what we can all do to be heroes going forward.

To quote Neil Jacobstein, “AI expands the range of the possible.” So keep doing it!

Develop the new ideas that will take Intelligent Assistants and other AI apps from the S-curve we are riding now to the next S-curve. I don’t know what will be the dominant technology of that next S-curve, but I do look forward to finding out.

And finally, what I asked you to remember at the outset: for AI to benefit humankind it has to be deployed; for successful deployment, good AI ideas have to be integrated into the human context of actual use and into the IT context of organizations.

Thanks

For sharing their ideas on high-impact applications and for their observations on what (we thought) we knew then / what (we think) we know now.

Alain Rappaport	Ashok Goel
Bruce Buchanan	David Stracuzzi
Edward Feigenbaum	Eric Schoen
Haym Hirsh	Howard Shrobe
Markus Fromherz	Mehmet Goker
Neil Jacobstein	Nestor Rychtyckyj
Peter Yeh	Phil Klahr
Ramasamy Uthurusamy	Ted Senator

I have only one thing left to do. This presentation is based on the ideas of many people. In the main, I have acted as a reporter.

Here they are.

Special Thanks

To Bruce Buchanan and Ed Feigenbaum for years of guidance ... and patience.

To Eric Schoen and Joshua Eckroth for their collaboration and for the IAAI content analytics.

To Jon Glick, AITopics pioneer.

To Raj Reddy, the Godfather of IAAI.

And to my family and friends for foundational contributions.

And special thanks are due these people.

Thank you very much for your attention ... and your accomplishments!

References

- Björnfors, Carl and Mattias Ellison. Enterprise Search and Findability Survey 2014. http://www2.findwise.com/findabilitysurvey2014
- Cheetham, W.E. Tenth Anniversary of the Plastics Color Formulation Tool. Presented at the *Sixteenth IAAI Conference (IAAI-2004)*, San Jose, California, July, 2004.
- Chun, A.H.W., T.Y.Y. Suen. Engineering Works Scheduling for Hong Kong's Rail Network. Presented at the *Twenty-Sixth IAAI Conference (IAAI-2014)*, Québec City, Québec, Canada, July, 2014. (Also see: "The AI boss that deploys Hong Kong's subway engineers", TECHNOLOGY NEWS, 2 July 2014. https://www.newscientist.com/article/mg22329764.000-the-ai-boss-that-deploys-hong-kongs-subway-engineers .)
- Dzierzanowski, J.M., K.R. Chrisman, G. J. MacKinnon, and P. Klahr. The Authorizer's Assistant: A Knowledge-Based Credit Authorization System for American Express. Presented at the *First IAAI Conference (IAAI-89)*, Stanford, CA, May, 1989.
- Feigenbaum, Edward A. Tiger in a Cage: The Applications of Knowledge-based Systems. Invited talk, presented at the AAAI-93 / IAAI-93 Conference, Washington, DC, July, 1993.

References

- Ferrucci, David, Eric Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya A. Kalyanpur, Adam Lally, J. William Murdock, Eric Nyberg, John Prager, Nico Schlaefer, and Chris Welty. Building Watson: An Overview of the DeepQA Project. AI Magazine, Vol 31, No. 3, Fall 2010, pp. 59-79.
- Katerina Frantzi, Sophia Ananiadou, Hideki Mima. Automatic Recognition of Multi-Word Terms: the C-value/NC-value Method. International Journal on Digital Libraries 08/2000; 3(2):115-130. DOI: 10.1007/3-540-49653-X_35.
- Foster, R. *Innovation: The Attacker's Advantage*, Summit Books, New York. 1986.
- Genovese, Yvonne (2012). Information and the Nexus of Forces: Delivering and Analyzing Data. Gartner Report ID: G00234835 (26 June 2012).
- Geun-Sik Jo, Kyeong-Jin Oh, Inay Ha, Kee-Sung Lee, Myung-Duk Hong, Ulrich Neumann, Suya You. A Unified Framework for Augmented Reality and Knowledge-Based Systems in Maintaining Aircraft. Presented at the *Twenty-Sixth IAAI Conference (IAAI-14)*, Québec City, Québec, CA, July, 2014.
- Halevy, A.Y., P. Norvig, F. Pereira. (2009). The Unreasonable Effectiveness of Data. IEEE Expert / IEEE Intelligent Systems - EXPERT , vol. 24, no. 2, pp. 8-12.

References

- Hammond, Kris. How artificial intelligence expains analytics. *Computerworld*, June 19, 2015. http://www.computerworld.com/article/2936142/emerging-technology/how-artificial-intelligence-explains-analytics.html
- Horvitz, Eric. One Hundred Year Study on Artificial Intelligence: Reflections and Framing. 2014. Unpublished memorandum. https://ai100.stanford.edu/reflections-and-framing
- Jacobstein, Neil. Innovative Applications of Early Stage AI. Presented at the 2007 Singularity Summit, San Francisco, CA, September, 2007. http://foresight.org/cms/events_library/267
- Kahneman, Daniel. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, 2011.
- Kirkland, J.D., T. E. Senator, J. J. Hayden, T. Dybala, H.G. Goldberg, P. Shyr. (1999). The NASD Regulation Advanced-Detection System (ADS). *AI Magazine* 20 (1): 55-68.
- Lohr, Steve. If Algorithms Know All, How Much Should Humans Help? *New York Times*, April 6, 2015. http://www.nytimes.com/2015/04/07/upshot/if-algorithms-know-all-how-much-should-humans-help.html

References

- Markoff, John. *Machines of Loving Grace: The Quest for Common Ground between Humans and Robots*. First edition. Ecco, an imprint of HarperCollins Publishers, 2015.
- Muratore, J.F., T. A. Heindel, T.i B. Murphy, A. N. Rasmussen, and R. Z. McFarland. Applications of Artificial Intelligence to Space Shuttle Mission Control. Presented at the *First IAAI Conference (IAAI-89)*, Stanford, CA, May, 1989.
- Preshing, Jeff. A Look Back at Single-Threaded CPU Performance. *Preshing on Programming Blog*, February 8, 2012. http://preshing.com/20120208/a-look-back-at-single-threaded-cpu-performance/
- Rychtyckyj, N. DLMS: Ten Years of AI for Vehicle Assembly Process Planning. Presented at the *Eleventh IAAI Conference (IAAI-99)*, Orlando, Florida, July, 1999.
- Sandholm, Tuomas. Expressive Commerce and Its Application to Sourcing: How We Conducted $35 Billion of Generalized Combinatorial Auctions. *AI Magazine*, Vol 28, No 3, Fall 2007, pp. 45-58.
- Shrobe, Howard E. The Innovative Applications of Artificial Intelligence Conference: Past and Future. *AI Magazine*, Vol 17, No 4, Winter 1996, pp. 15-20.

References

- Shrobe, Howard E. What Does the Future Hold? *AI Magazine*, Vol 21, No 4, Winter 2000, pp. 41-57.
- R.G. Smith. On the Development of Commercial Expert Systems. *AI Magazine*, 5(3):61-73, Fall 1984. Reprinted in *Readings from AI Magazine,*Volumes 1-5, 1980-1985, R. Engelmore, editor, American Association for Artificial Intelligence, 1988.
- Smith, Reid G. and Adam Farquhar. The Road Ahead for Knowledge Management: An AI Perspective. *AI Magazine*, 21(4):17-40, Winter 2000.
- Thompson, E.D., E. Frolich, J. C. Bellows, B. E. Bassford, E. J. Skiko, M. S. Fox. Process Diagnosis System (PDS) - A 30 Year History. Presented at the *Twenty-Seventh IAAI Conference (IAAI-15)*, Austin, TX, January, 2015.
- Thrun, Sebastian. A Personal Account of the Development of Stanley, the Robot That Won the DARPA Grand Challenge. *AI Magazine*. Vol 27, No 4, Winter 2006, pp. 69-82.
- Urmson, C., C. Baker, J. Dolan, P. Rybski, B. Salesky, W. Whittaker, D. Ferguson, M. Darms. Autonomous Driving in Traffic: Boss and the Urban Challenge. Presented at the *Twentieth IAAI Conference (IAAI-08)*, Chicago, IL, July, 2008.

References

- Weng, G.K., G. K.Seng, L.C. Whye, T.B. Hwa, T.E. Kiang. Scheduling of Port of Singapore Authority: A Total Approach. Presented at the *Seventh IAAI Conference (IAAI-95)*, Montreal, Quebec, Canada, August, 1995.
- Winston, H.A., R. T. Clark, G. Buchina. AGETS MBR An Application of Model-Based Reasoning to Gas Turbine Diagnostics. Presented at the *Seventh IAAI Conference (IAAI-95)*, Montreal, Quebec, Canada, August, 1995.
- Wu, J., K.M. Williams, H-H. Chen, M. Khabsa, C. Caragea, S. Tuarob, A. G. Ororbia, D. Jordan, P. Mitra, C. L. Giles. CiteSeerX: AI in a Digital Library Search Engine. Presented at the *Twenty-Fifth IAAI Conference (IAAI-13)*, Bellevue, WA, August, 2013.