

COMPUTING  
**edge**

**ARTIFICIAL  
INTELLIGENCE**

Also in this issue:

- > Risks
- > Moving to Autonomous and Self-Migrating Containers for Cloud Applications

APRIL 2017

[www.computer.org](http://www.computer.org)

 **IEEE**

IEEE  computer society



## Focus on Your Job Search

**IEEE Computer Society Jobs** helps you easily find a new job in IT, software development, computer engineering, research, programming, architecture, cloud computing, consulting, databases, and many other computer-related areas.

**New feature:** Find jobs recommending or requiring the IEEE CS CSDA or CSDP certifications!

Visit [www.computer.org/jobs](http://www.computer.org/jobs) to search technical job openings, plus internships, from employers worldwide.

<http://www.computer.org/jobs>

IEEE  computer society | **JOBS**



The IEEE Computer Society is a partner in the AIP Career Network, a collection of online job sites for scientists, engineers, and computing professionals. Other partners include Physics Today, the American Association of Physicists in Medicine (AAPM), American Association of Physics Teachers (AAPT), American Physical Society (APS), AVS Science and Technology, and the Society of Physics Students (SPS) and Sigma Pi Sigma.



STAFF

**Editor**  
Lee Garber

**Contributing Staff**  
Christine Anthony, Brian Brannon, Lori Cameron, Cathy Martin, Chris Nelson, Meghan O'Dell, Dennis Taylor, Rebecca Torres, Bonnie Wylie

**Production & Design**  
Carmen Flores-Garvey, Monette Velasco, Jennie Zhu-Mai, Mark Bartosik

**Acting Editorial Content Manager**  
Carrie Clark

**Senior Manager, Editorial Services**  
Robin Baldwin

**Director, Products and Services**  
Evan Butterfield

**Senior Advertising Coordinator**  
Debbie Sims

**Circulation:** ComputingEdge (ISSN 2469-7087) is published monthly by the IEEE Computer Society. IEEE Headquarters, Three Park Avenue, 17th Floor, New York, NY 10016-5997; IEEE Computer Society Publications Office, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720; voice +1 714 821 8380; fax +1 714 821 4010; IEEE Computer Society Headquarters, 2001 L Street NW, Suite 700, Washington, DC 20036.

**Postmaster:** Send address changes to ComputingEdge-IEEE Membership Processing Dept., 445 Hoes Lane, Piscataway, NJ 08855. Periodicals Postage Paid at New York, New York, and at additional mailing offices. Printed in USA.

**Editorial:** Unless otherwise stated, bylined articles, as well as product and service descriptions, reflect the author's or firm's opinion. Inclusion in ComputingEdge does not necessarily constitute endorsement by the IEEE or the Computer Society. All submissions are subject to editing for style, clarity, and space.

**Reuse Rights and Reprint Permissions:** Educational or personal use of this material is permitted without fee, provided such use: 1) is not made for profit; 2) includes this notice and a full citation to the original work on the first page of the copy; and 3) does not imply IEEE endorsement of any third-party products or services. Authors and their companies are permitted to post the accepted version of IEEE-copyrighted material on their own Web servers without permission, provided that the IEEE copyright notice and a full citation to the original work appear on the first screen of the posted copy. An accepted manuscript is a version which has been revised by the author to incorporate review suggestions, but not the published version with copy-editing, proofreading, and formatting added by IEEE. For more information, please go to: [http://www.ieee.org/publications\\_standards/publications/rights/paperversionpolicy.html](http://www.ieee.org/publications_standards/publications/rights/paperversionpolicy.html). Permission to reprint/republish this material for commercial, advertising, or promotional purposes or for creating new collective works for resale or redistribution must be obtained from IEEE by writing to the IEEE Intellectual Property Rights Office, 445 Hoes Lane, Piscataway, NJ 08854-4141 or [pubs-permissions@ieee.org](mailto:pubs-permissions@ieee.org). Copyright © 2017 IEEE. All rights reserved.

**Abstracting and Library Use:** Abstracting is permitted with credit to the source. Libraries are permitted to photocopy for private use of patrons, provided the per-copy fee indicated in the code at the bottom of the first page is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

**Unsubscribe:** If you no longer wish to receive this ComputingEdge mailing, please email IEEE Computer Society Customer Service at [help@computer.org](mailto:help@computer.org) and type "unsubscribe ComputingEdge" in your subject line.

IEEE prohibits discrimination, harassment, and bullying. For more information, visit [www.ieee.org/web/aboutus/whatis/policies/p9-26.html](http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html).

## IEEE Computer Society Magazine Editors in Chief

**Computer**

Sumi Helal, *University of Florida*

**IEEE Software**

Diomidis Spinellis, *Athens University of Economics and Business*

**IEEE Internet Computing**

M. Brian Blake, *University of Miami*

**IT Professional**

San Murugesan, *BRITE Professional Services*

**IEEE Security & Privacy**

Ahmad-Reza Sadeghi, *Technical University of Darmstadt*

**IEEE Micro**

Lieven Eeckhout, *Ghent University*

**IEEE Computer Graphics and Applications**

L. Miguel Encarnaçao, *ACT, Inc.*

**IEEE Pervasive Computing**

Maria Ebling, *IBM T.J. Watson Research Center*

**Computing in Science & Engineering**

Jim X. Chen, *George Mason University*

**IEEE Intelligent Systems**

V.S. Subrahmanian, *University of Maryland*

**IEEE MultiMedia**

Yong Rui, *Lenovo Research and Technology*

**IEEE Annals of the History of Computing**

Nathan Ensmenger, *Indiana University Bloomington*

**IEEE Cloud Computing**

Mazin Yousif, *T-Systems International*

# COMPUTING edge



14

Artificial Intelligence as an Effective Classroom Assistant

26

There Is No AI Without IA

36

Masterminds of Artificial Intelligence: Marvin Minsky and Seymour Papert



# 39

Risks

- 7 Editor's Note:  
A Close Look at Artificial Intelligence
- 8 Machine Learning  
PANOS LOURIDAS AND CHRISTOF EBERT
- 14 Artificial Intelligence as an Effective Classroom  
Assistant  
BENEDICT DU BOULAY
- 20 TJBOT and Zooids: The Connection between  
Pervasive Computing and AI  
MARIA R. EBLING
- 23 Engineering the New Boundaries of AI  
AMIR BANIFATEMI AND JEAN-LUC GAUDIOT
- 26 There Is No AI Without IA  
SETH EARLEY
- 33 Unreliable Machine Learning  
VINTON G. CERF
- 36 Masterminds of Artificial Intelligence: Marvin  
Minsky and Seymour Papert  
GEORGE STRAWN AND CANDACE STRAWN
- 39 Risks  
MATTHEW TURK
- 42 Moving to Autonomous and Self-Migrating  
Containers for Cloud Applications  
DAVID S. LINTHICUM

## Departments

- 4 Magazine Roundup
- 46 Computing Careers:  
Careers in Artificial Intelligence
- 48 Career Opportunities

Subscribe to **ComputingEdge** for free at  
[www.computer.org/computingedge](http://www.computer.org/computingedge).

# Magazine Roundup

The IEEE Computer Society's lineup of 13 peer-reviewed technical magazines covers cutting-edge topics ranging from software design and computer graphics to Internet computing and security, from scientific applications and machine intelligence to cloud migration and microchip design. Here are highlights from recent issues.

## *Computer*

Our society increasingly relies on the successful marriage of digital and physical systems to perform advanced automation and control tasks. Thus, engineering these **cyber-physical systems** to

ensure their efficiency, security, and dependability is an important research area. Meanwhile, as essential, life-sustaining systems increasingly connect in cyberspace, we must figure out how to **identify, assess, anticipate, and manage risk**. *Computer's* April 2017 issue has special sections on these topics.

## *IEEE Software*

With global software engineering becoming standard practice, today's software engineering students will become tomorrow's global software engineers. The education systems underpinning the profession will thus need to change accordingly.

Current approaches to teaching software engineering are outdated and don't address today's conditions. In "How Best to Teach Global Software Engineering? Educators Are Divided," from *IEEE Software's* January/February 2017 issue, pioneering educators discuss how they inject realism into **global software engineering education**.

## *IEEE Internet Computing*

**Fog computing** could alleviate many of the Internet of Things' challenges by, for example, reducing latency and enhancing security. *IEEE Internet Computing's* March/April 2017 special issue explores the technology's potential to both form distributed and virtualized platforms that support computation-intensive tasks and distribute advanced computing, storage, networking, and management services to the network edge. The issue also looks at the obstacles fog computing faces.

## *Computing in Science & Engineering*

*CiSE's* March/April 2017 special issue provides a mix of perspectives for understanding more deeply where **Moore's law** stands today and the prospects for continued exponential performance growth in the upcoming computing era.

## *IEEE Security & Privacy*

*IEEE S&P's* January/February 2017 special issue looks at how **the latest security and privacy research** is moving the field forward. Articles in the issue address topics such as the use of containerization to secure mobile devices, cloud data-auditing techniques, and a novel security-development lifecycle model.

## *IEEE Cloud Computing*

*IEEE Cloud Computing's* January/February 2017 issue includes articles on subjects such as **combating distributed denial-of-service attacks in the cloud**, secure data sharing and searching at the edge of the cloud-assisted Internet of Things, and trusting cloud service providers.

## *IEEE Computer Graphics and Applications*

Articles from *CG&A's* March/April 2017 special issue on **information visualization** examine topics such as layered graph drawing for visualizing evaluation structures, the visualization of rank time series of top-viewed Wikipedia pages, and spatial analytic interfaces.

## *IEEE Intelligent Systems*

"An Open Semantic Framework for the Industrial Internet of Things," from *IEEE Intelligent Systems' January/February 2017 issue*, introduces the **Open Semantic Framework (OSF)**, which lets users create, deploy, and manage semantic applications. The OSF represents an important step toward the ability to integrate various knowledge models.

## *IEEE MultiMedia*

The authors of "The Continuing Reinvention of Content-Based Retrieval: Multimedia Is Not Dead," from *IEEE MultiMedia's January-March 2017 special issue*, review how **content-based retrieval** has evolved over the years and discuss the technology's future goals and research directions.

## *IEEE Annals of the History of Computing*

As part of a broader look at **early intrusion-detection systems**, "The March of IDES: Early History of Intrusion-Detection Expert Systems," from *IEEE Annals' October-December 2016 issue*, focuses on the first such application, the Intrusion Detection Expert System (IDES), which SRI International developed in the 1980s.

## *IEEE Pervasive Computing*

The authors of "**The Future of Pervasive Health**," from *IEEE Pervasive Computing's January-March 2017 issue*, discuss five important matters that must be addressed

before the pervasive-health field can advance: technological challenges and opportunities, adoption and adherence, open data, ethical issues, and education.

## *IT Professional*

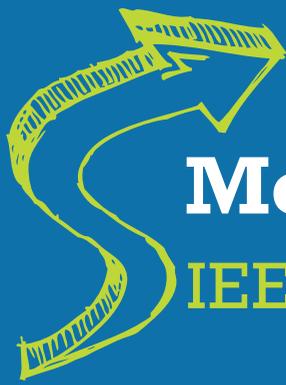
To achieve digital transformation, organizations must control the avalanche of data generated by supplier, customer, and employee interactions. This complex data environment has led enterprises to realize that new leadership—in the form of a **chief data officer (CDO)**—is needed to oversee this task and to take responsibility for extracting value from data. "The Evolving Role of the CDO," from *IT Pro's January/February 2017 issue*, examines this relatively new position.

## *IEEE Micro*

*IEEE Micro's* January/February 2017 special issue on **cognitive architectures** includes articles on spacetime computation and the neocortex, low-power automatic speech recognition via a mobile GPU and Viterbi accelerator, efficient situational scheduling of graph workloads on single-chip multicores and GPUs, and graph-analytics accelerators for cognitive systems.

## *Computing Now*

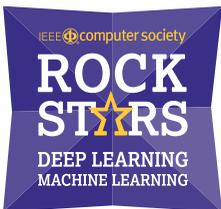
The Computing Now website ([computingnow.computer.org](http://computingnow.computer.org)) features **up-to-the-minute computing news** and blogs, along with articles ranging from peer-reviewed research to opinion pieces by industry leaders. ●



# Move Your Career Forward

## IEEE Computer Society Membership

### Explore These Artificial Intelligence Resources



#### **Rock Stars of Deep Learning/Machine Learning**

**12 September 2017**

**Mountain View, CA, USA**

Join us in Silicon Valley with Fortune 100 C-suite and upper management executives and others to learn about deep learning and machine learning opportunities for 2017. Hear insights on how they use analytics, algorithms, and data structures, and learn how to trust semisupervised or unsupervised analysis and how to contract for or build your deep learning talent pool.

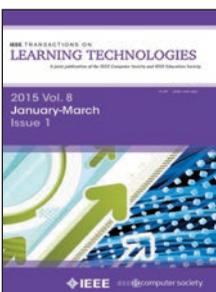


#### **Rock Stars of Smart Robo Technologies**

**17 October 2017**

**Dallas, TX, USA**

This conference will explore how companies are using smart robo systems for financial services, insurance, healthcare, manufacturing, publishing, and marketing. Join us in the heart of Texas, as executive speakers will discuss their current implementations and expectations for future applications. They'll address key issues, such as integration strategies and challenges; ROI; impact on the workforce; accuracy; and security.



#### ***IEEE Transactions on Learning Technologies***

The quarterly journal *TLT* covers research on such topics as innovative online learning systems, intelligent tutors, educational software applications and games, and simulation systems for education and training.

**FOR DIRECT LINKS  
TO THESE RESOURCES, VISIT**

[www.computer.org/edge-resources](http://www.computer.org/edge-resources)

IEEE  computer society

*The Community for Technology Leaders*

# A Close Look at Artificial Intelligence

Long the subject of science fiction, artificial intelligence (AI) began as a formal research discipline at a 1956 Dartmouth College workshop. Since then, optimism about artificial intelligence has risen and fallen, based in part on the limitations of AI itself and the technologies it uses. Now, AI is being employed extensively in business applications and consumer products, giving the technology an economic foundation.

This *ComputingEdge* issue examines today's important AI-related developments.

*IEEE Software's* "Machine Learning" presents a brief overview of machine-learning technologies, with a case study involving code analysis.

"Artificial Intelligence as an Effective Classroom Assistant," from *IEEE Intelligent Systems*, reviews metareviews and meta-analyses to make the case for blended learning, in which a teacher offloads some work to systems that use AI and cognitive science to help with the process.

The author of "TJBot and Zooids: The Connection between Pervasive Computing and AI," from *IEEE Pervasive Computing*, looks at a robot and a robot kit to show how AI and pervasive computing could reshape the way that humans interact with computers.

*Computer's* "Engineering the New Boundaries of AI" discusses the XPRIZE Foundation's launching of the IBM Watson AI XPRIZE, in which

interdisciplinary teams advance AI technologies while competing for the grand prize.

The author of "There Is No AI Without IA," from *IT Professional*, argues that artificial intelligence needs high-quality, structured data—in part via effective information architectures (IAs)—to be useful to organizations and their customers.

In *IEEE Internet Computing's* "Unreliable Machine Learning," Internet pioneer Vinton Cerf speculates about whether neural networks might experience problems as users ask them to work with more complex and diverse parameters and with larger decisions and responsibilities at stake.

*IT Professional's* "Masterminds of Artificial Intelligence: Marvin Minsky and Seymour Papert" highlights the contributions of two AI pioneers.

*ComputingEdge* articles on topics other than AI include the following:

- *Computing in Science & Engineering's* "Risks" focuses on scientific computing's potential hazards.
- The author of "Moving to Autonomous and Self-Migrating Containers for Cloud Applications," from *IEEE Cloud Computing*, says containers should be able to migrate from one type of public cloud to another in automated and autonomous ways. He then describes the manner in which this could be done. 📍



# Machine Learning

Panos Louridas and Christof Ebert

Machine learning is the major success factor in the ongoing digital transformation across industries. Startups and behemoths alike announce new products that will learn to perform tasks that previously only humans could do, and perform those tasks better, faster, and more intelligently. But how do they do it? What does it mean for IT developers and software engineers? Here, Panos Louridas and I present a brief overview of machine-learning technologies, with a concrete case study from code analysis. I look forward to hearing from both readers and prospective column authors. —Christof Ebert

**MACHINE LEARNING ISN'T** new; it has been around at least since the 1970s, when the first related algorithms appeared. What has changed is that the explosion in computing power has allowed us to use machine learning to tackle ever-more-complex problems, while the explosion of data being captured and stored has allowed us to apply machine learning to an ever-expanding range of domains.

Machine learning is used in different domains. Here are a few examples:

- security heuristics that distill attack patterns to protect, for instance, ports or networks;
- image analysis to identify distinct forms and shapes, such as for medical analyses or face and fingerprint recognition;
- deep learning to generate rules for data analytics and big data handling, such as are used in marketing and sales promotions;
- object recognition and predictions from combined video streams and

multisensor fusion for autonomous driving; and

- pattern recognition to analyze code for weaknesses such as criticality and code smells (for a related case study, see the sidebar).

The general idea behind most machine learning is that a computer learns to perform a task by studying a training set of examples. The computer (or system of distributed or embedded computers and controllers) then performs the same task with data it hasn't encountered before.

## Learning Strategies

Machine learning employs the following two strategies (see Figure 1).

### Supervised Learning

In *supervised learning*, the training set contains data and the correct output of the task with that data. This is like giving a student a set of problems and their solutions and telling that student to figure out





## CASE STUDY: MACHINE LEARNING FOR CODE ANALYSIS

Machine learning has many practical applications. As befits this magazine, we'll present an example that shows how machine learning can manage quality risks and improve quality assurance productivity.

### A CRITICALITY ASSESSMENT TOOL

Project managers and product owners often wonder when would be the right release point and how to assess the criticality of the code to be delivered. Static-analysis tools exist but give numerous warnings that might be difficult to link to actual weak spots. At Vector Consulting Services, we offer clients a machine-learning-based criticality assessment tool for software release management. Criteria include hard factors from static code analysis, such as cyclomatic complexity, the degree of reuse, and the history of defects in preceding versions and variants. We also use soft factors such as designer competences, experience with similar projects, and architectural decisions that might incur technical debt.

Using these criteria, the machine-learning tool builds a criticality prediction model. On the basis of a ranked list of the criticality of the modules used in a build, developers can apply different mechanisms to improve quality—refactoring, redesign, thorough static analysis, and unit testing with increased coverage schemes.

Instead of predicting the number of defects or changes (algorithmic relationships), the tool considers assignments to classes (for example, “defect prone”). The training and test data come from finished projects that had been under configuration control since coding started. After the machine-learning process, the data used operationally comes from active projects.

To achieve feedback to improve predictions, this approach is integrated throughout development (requirements, design, code, system test, and deployment).

### STEP-BY-STEP CRITICALITY PREDICTION

Figure A illustrates criticality classification and validation, which consists of eight steps.

**Step 1:** For the finished projects, provide a list of all the modules used for learning from the configuration system.

**Step 2:** Provide a defect list for each learning module. For high-ranking defects, you might add a root-cause analysis that allows for a Pareto-based mitigation list.

**Step 3:** Provide a change history classification (that is, the number of compiles or deliveries) for each learning module.

**Step 4:** With static code analysis, assemble for each learning module a complexity classification such as hot spots from code analysis.

**Step 5:** With the machine-learning system, construct an initial criticality list that takes into account the inputs from steps 2, 3, and 4, mapped to the list from step 1. Evaluate the criticality list's validity—for example, by screening for the identified critical modules, outliers, and potential misleading effects. Such screening aims to find undesired influences from the defect or change histories. The screening and ranking must primarily ensure the fewest possible type-I prediction errors. (In type-I errors, defect-prone components are misclassified as uncritical components.)

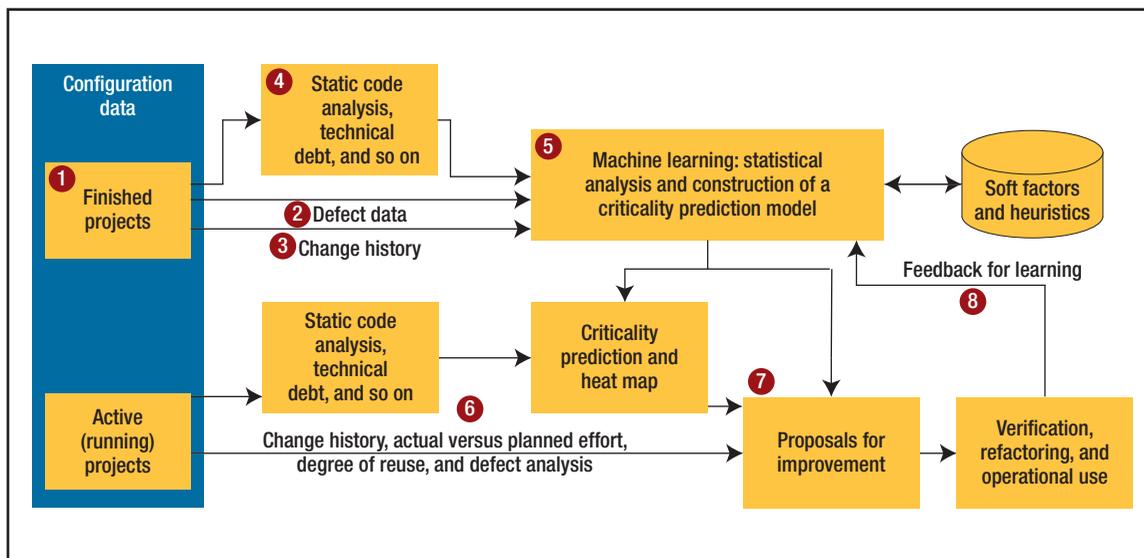
**Step 6:** For the current project, repeat steps 1 to 5 to get a predictive result on each new module's criticality. Then, present the rankings so that the developers can decide on further actions.

**Step 7:** Manually prepare suggestions based on the new modules ranked the most critical. Critical modules should at least undergo a flash review and subsequent refactoring, redesign, or rewriting—depending on their complexity, age, and reuse in other projects. Refactoring includes reducing size, improving modularity, balancing cohesion and coupling, and so on. For instance, apply thorough unit testing with 100 percent CO coverage (statement coverage) to the modules. Investigate the details of those modules' complexity measurements to determine the redesign approach. Typically, the different complexity measurements will indicate the approach to follow.

**Step 8:** After the new project is finished, validate and improve the prediction model on the basis of postmortem studies with all collected defects and the population of a “real” criticality list. Then, compare the actual defect ranking with the predicted ranking. Investigate the reasons for deviations, and tune the implemented automatic classification approaches. Improve the screening rules to ensure that type-II prediction errors will be reduced the next time. (In type-II errors, uncritical components are misclassified as defect-prone components.)

*continued on next page*

continued from previous page



**FIGURE A.** Machine learning for criticality prediction of source code as implemented in this case study. The numbers refer to the steps described in the sidebar text.

## THIS APPROACH'S EFFECTIVENESS

Criticality prediction doesn't aim to detect all defects. Instead, it aims to optimize resource allocation by focusing resources on areas with critical defects that would affect the delivered product's utility. We estimate the trade-off of the costs of applying complexity-based predictive quality models and the eventual code changes versus the improved quality, on the following basis:

- Effort—limited resources are assigned to the high-risk components.
- Effort—gray-box testing strategies are applied to only the high-risk components.
- Benefits—the risk assessment of changes is eased with code analysis on the basis of the affected or changed complexity.
- Benefits—fewer customers reported failures than in previous releases, and security and maintainability improved.

On the basis of the results from many of our client projects (and taking a conservative ratio of only 40 percent defects in critical components), we can calculate the

business case:

- On average, for each project, 20 percent of the modules were selected as the most critical (after coding).
- Those modules contained over 40 percent of all defects (up to release time).

In addition, we've determined that 60 percent of all defects theoretically can be detected until the end of unit testing. Also, defect correction with unit testing and static analysis costs 10 to 50 percent less than defect correction in subsequent testing activities. So, we calculate that developers can detect 24 percent of all defects early by investigating 20 percent of all modules more intensively, with over 10 percent reduced effort than with late defect correction. This yields at least a 20 percent cost reduction for defect correction.

The necessary tools, such as Coverity, Klocwork, Lattix, Structure 101, SonarX, and SourceMeter, are off the shelf and account for even less per project. These criticality analyses provide numerous other benefits, such as the removal of specific code-related risks and defects that otherwise are hard to identify (for example, security flaws).

how to solve other problems he or she will have to deal with in the future.

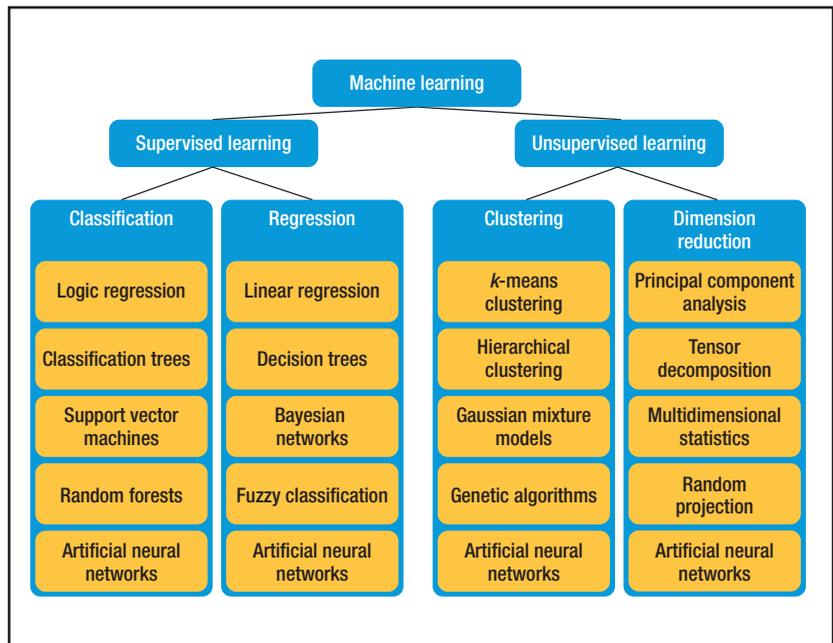
Supervised learning includes *classification* algorithms, which take as input a dataset and the class of each piece of data so that the computer can learn how to classify new data. For example, the input might be a set of past loan applications with an indication of which of them went bad. On the basis of this information, the computer classifies new loan applications. Classification can employ logic regression, classification trees, support vector machines, random forests, artificial neural networks (ANNs), or other algorithms. ANNs are a major topic on their own; we discuss them in more detail later.

*Regression* algorithms predict a value of an entity's attribute ("regression" here has a wider sense than merely statistical regression). Regression algorithms include linear regression, decision trees, Bayesian networks, fuzzy classification, and ANNs.

### Unsupervised Learning

In *unsupervised learning*, the training set contains data but no solutions; the computer must find the solutions on its own. This is like giving a student a set of patterns and asking him or her to figure out the underlying motifs that generated the patterns.

Unsupervised learning includes *clustering* algorithms, which take as input a dataset covering various dimensions and partition it into clusters satisfying certain criteria. A popular algorithm is *k*-means clustering, which aims to partition the dataset so that each observation lies closest to the mean of its cluster. Other clustering approaches include hierarchical clustering, Gaussian mixture models, genetic algorithms (in which the computer learns the



**FIGURE 1.** Machine-learning approaches. In machine learning, a computer first learns to perform a task by studying a training set of examples. The computer then performs the same task with data it hasn't encountered before.

best way for a task through artificial selection), and ANNs.

*Dimensionality reduction* algorithms take the initial dataset covering various dimensions and project the data to fewer dimensions. These fewer dimensions try to better capture the data's fundamental aspects. Dimensionality reduction algorithms include principal component analysis, tensor reduction, multidimensional statistics, random projection, and ANNs.

### Essential Tools

Machine learning's popularity has brought along a wealth of tools. Most of them are open source, so users can easily experiment with them and learn how to use them. Table 1 compares some popular machine-learning tools.

The numerical and statistical communities are divided into two

camp: one that prefers R and one that prefers Python. Of course, any absolute division makes no sense. For a field as wide as machine learning, no single tool will do. The best a software engineer can do is to become acquainted with many different tools and learn which one is the most appropriate for a given situation.

That said, R is more popular with people with a somewhat stronger statistical background. It has a superb collection of machine-learning and statistical-inference libraries. Chances are, if you find a fancy algorithm somewhere and want to try it on your data, an implementation in R exists for it. R boasts the ggplot2 visualization library, which can produce excellent graphs.

Python is more popular with people with a computer science background. Although not made specifically for machine learning or

TABLE 1

Some popular machine-learning tools.

	Tool				
	Python	R	Spark	Matlab	TensorFlow
License	Open source	Open source	Open source	Proprietary	Open source
Distributed	No	No	Yes	No	No
Visualization	Yes	Yes	No	Yes	No
Neural nets	Yes	Yes	Multilayer perceptron classifier	Yes	Yes
Supported languages	Python	R	Scala, Java, Python, and R	Matlab	Python and C++
Variety of machine-learning models	High	High	Medium	High	Low
Suitability as a general-purpose tool	High	Medium	Medium	High	Low
Maturity	High	Very high	Medium	Very high	Low

statistics, Python has extensive libraries for numerical computing (NumPy), scientific computing (SciPy), statistics (StatsModels), and machine learning (scikit-learn). These are largely wrappers of C code, so you get Python's convenience with C's speed.

Although there are fewer machine-learning libraries for Python than there are for R, many programmers find working with Python easier. They might already know the language or find it easier to learn than R. They also find Python convenient for preprocessing data: reading it from various sources, cleaning it, and bringing it to the required formats. For visualization, Python relies on matplotlib. You can do pretty much everything on matplotlib, but you might discover you have to put in some effort. The seaborn library is built on top of it, letting you produce elegant visualizations with little code.

In general, R and Python work when the dataset fits in the com-

puter's main memory. If that's not possible, you must use a distributed platform. The most well-known is Hadoop, but Hadoop isn't the most convenient for machine learning. Making even simple algorithms run on it can be a struggle.

So, many people prefer to work at the higher level of abstraction that Spark offers. Spark leverages Hadoop but looks like a scripting environment. You can interact with it using Scala, Java, Python, or R. Spark has a machine-learning library that implements key algorithms, so for many purposes you don't need to implement anything yourself.

H2O is a relatively newer entrant in the machine-learning scene. It's a platform for descriptive and predictive analytics that uses Hadoop and Spark; you can use it with R and Python. It implements supervised- and unsupervised-learning algorithms and a Web interface through which you can organize your workflow.

A promising development is the

Julia programming language for technical computing, which aims at top performance. Because Julia is new, it doesn't have nearly as many libraries as Python or R. Yet, thanks to its impressive speed, its popularity might grow.

Strong commercial players include Matlab and SAS, which both have a distinguished history. Matlab has long offered solid tools for numerical computation, to which it has added machine-learning algorithms and implementations. For engineers familiar with Matlab, it might be a natural fit. SAS is a software suite for advanced statistical analysis; it also has added machine-learning capabilities and is popular for business intelligence tasks.

### ANNs and Deep Learning

Cynics might roll their eyes, arguing that ANNs' resurgence is déjà vu. It's true; ANNs' fundamental components have been around for about half a century. However, it's also true

that you can now architect them in new ways. ANNs can be used across the spectrum of machine learning: classification, regression, clustering, and dimensionality reduction.

Innovations in ANN architectures and the availability of cheap computing resources to run ANNs has brought about the burgeoning of *deep learning*—using big ANNs to perform machine learning. Over the last few years, deep learning has chalked up headline-grabbing successes by beating humans in *Jeopardy!* and Go, learning to play arcade games, showing an uncanny capability to recognize images, performing automatic translation, and so on. Deep learning is particularly good at general tasks requiring the elicitation of higher-level, abstract concepts from the input data, which is what the many layers of an ANN excel at.

Deep learning is usually implemented through matrices, so working with it requires efficient matrix operations and manipulation. Usually the implementations are in C or C++, but designing ANNs at that level is unwieldy. Python programmers can use the Theano library to define ANNs, which are compiled to C code that's then compiled to machine language. Recently, Google released as open source its TensorFlow library for working with ANNs. You can interact with TensorFlow through a Python API. A C++ API is also available; although not as easy to use, it might give some performance benefits.

Before jumping on the deep-learning bandwagon, keep in mind that all machine-learning approaches lie on a spectrum based on the ease of interpreting their results. For example, classification trees produce rules that classify data. By reading

those rules, you can easily understand how a classification tree classifies data. ANNs don't produce anything their users can interpret. An ANN that classifies images doesn't produce any rules; the network itself embodies everything it has learned about image classification.

**M**any machine-learning books have a practical slant, aiming to introduce machine learning on a particular platform. As technologies quickly evolve, it's better to focus on getting a solid grasp of the fundamentals. After all, using a machine-learning platform isn't difficult; knowing when to use a particular algorithm and how to use it well requires quite a bit of background knowledge. Here are four popular books:

- T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining,*

*Inference, and Prediction*, 2nd ed., Springer, 2009.

- C.M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006.
- K.P. Murphy, *Machine Learning: A Probabilistic Perspective*, MIT Press, 2012.
- E. Alpaydın, *Introduction to Machine Learning*, 3rd ed., MIT Press, 2014. 

**PANOS LOURIDAS** is an associate professor teaching algorithms and software at the Athens University of Economics and Business. He's also an active developer. Contact him at [louridas@aub.gr](mailto:louridas@aub.gr).

**CHRISTOF EBERT** is the managing director of Vector Consulting Services. He is on the *IEEE Software* editorial board and teaches at the University of Stuttgart and the Sorbonne in Paris. Contact him at [christof.ebert@vector.com](mailto:christof.ebert@vector.com).

*This article originally appeared in IEEE Software, vol. 33, no. 5, 2016.*

## Take the CS Library wherever you go!



IEEE Computer Society magazines and Transactions are now available to subscribers in the portable ePub format.

Just download the articles from the IEEE Computer Society Digital Library, and you can read them on any device that supports ePub. For more information, including a list of compatible devices, visit

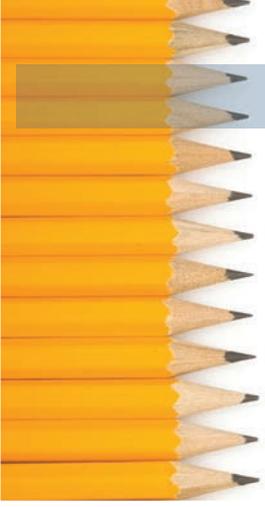
[www.computer.org/epub](http://www.computer.org/epub)



IEEE



IEEE computer society



# Artificial Intelligence as an Effective Classroom Assistant

**Benedict du Boulay**, *University of Sussex*

**T**he field of artificial intelligence in education (AIED) has been in existence for about 40 years and has operated under various other names, the most common of which is intelligent tutoring systems (ITSs). The field was initially brought to wider attention by papers in a special issue of the *International Journal of Man-Machine Studies*,<sup>1</sup> in a book based on that special issue,<sup>2</sup> and in AI books of the era.<sup>3</sup> It continues to use techniques from AI and cognitive science to attempt to understand the nature of learning and teaching and to build systems to assist learners to master new skills or understand new concepts in ways that mimic the actions of a skilled human tutor working one-on-one with the learner. That is, such systems attempt to adapt the way they teach to the learner's knowledge, skill, and preferred ways of learning, and to consider the learners' affective trajectory as they deal with the expected setbacks and impasses of mastering new material. There is clearly some overlap with other uses of computing technology in education, although the commitment to individual adaptation through modeling different parts of the educational process is key.

For such systems to adapt to the learner and provide a personalized experience, a typical conceptual architecture has evolved. This consists of

- a model of the domain being learned, so that the system can reason about and judge whether a student's answer or a problem-solving step is appropriate;
- a model of the learner's current understanding or skill level, so that tasks of appropriate complexity can be posed;
- a model of pedagogy, so that the system can make sensible tutorial moves, such as providing effective feedback or adjusting the nature of the next task; and
- one or more interfaces through which the system and the learner can communicate to explore and learn about the domain in question.

Over the years, many systems using various pedagogical techniques and topics have been built and evaluated. To illustrate the scope of the work, I discuss four diverse systems, which range from classic teaching in a formal subject and a procedural skill, to learning by creating externalized forms of knowledge for a highly conceptual learning task, to rich, natural user interaction via speech for learning complex, culture-laden skills.

The first AIED example is a system to help learners understand basic algebra by being given problems and then provided with step-by-step feedback and guidance for their solution.<sup>4</sup> The second example is a system that helps learners gain a conceptual understanding of river ecosystems by building a concept map of that domain, as if for another learner, and having that simulated other learner take tests on the concept map's adequacy.<sup>5</sup> The third example is a system that helps military personnel learn to speak Arabic and understand the social and cultural norms needed to interact with people in the country in which they are operating.<sup>6</sup>

The fourth example illustrates the increasing importance of the interface in AIED systems and their use in informal learning environments, such as museums, as well as formal ones. Figure 1 shows Coach Mike, a pedagogical agent designed to help children visiting a museum learn about

## Answer-, Step-, and Substep-Based Tutoring

Imagine that a student must solve the following equation:

$$2(14 - x) = 23 + 3x.$$

An *answer-based system* would expect the student to do all the work offline and then provide the answer  $x = 1$ . If asked for a hint, the tutor can suggest broad ways of going about the problem, such as to collect all the terms in  $x$  on one side of the equation, but the tutor has no way of knowing that this advice is being followed. If the answer provided is wrong—for example,  $x = 1.25$ —the tutor might be able to hypothesize that the student multiplied out the bracket incorrectly, but if the answer provided is, say,  $x = 14$ , it probably will not offer much in the way of specific help.

In a *step-based system*, the student might be invited to multiply out the bracket expression as a first step, and thus will give  $28 - 2x$  as the answer to that step. If the student requests a hint or gives a wrong answer to this step, the tutor can give help about how to work that step. Once the step is completed correctly, the tutor would invite an answer to the next step, such as reordering terms in the equation, and then on through further steps to the final answer.

In a *substep-based system*, there might be a remedial dialogue at a finer level than an individual step, for instance, about what expressions such as  $2x$  or  $3x$  mean, if that seems warranted by the request for a hint or by a wrong step answer.

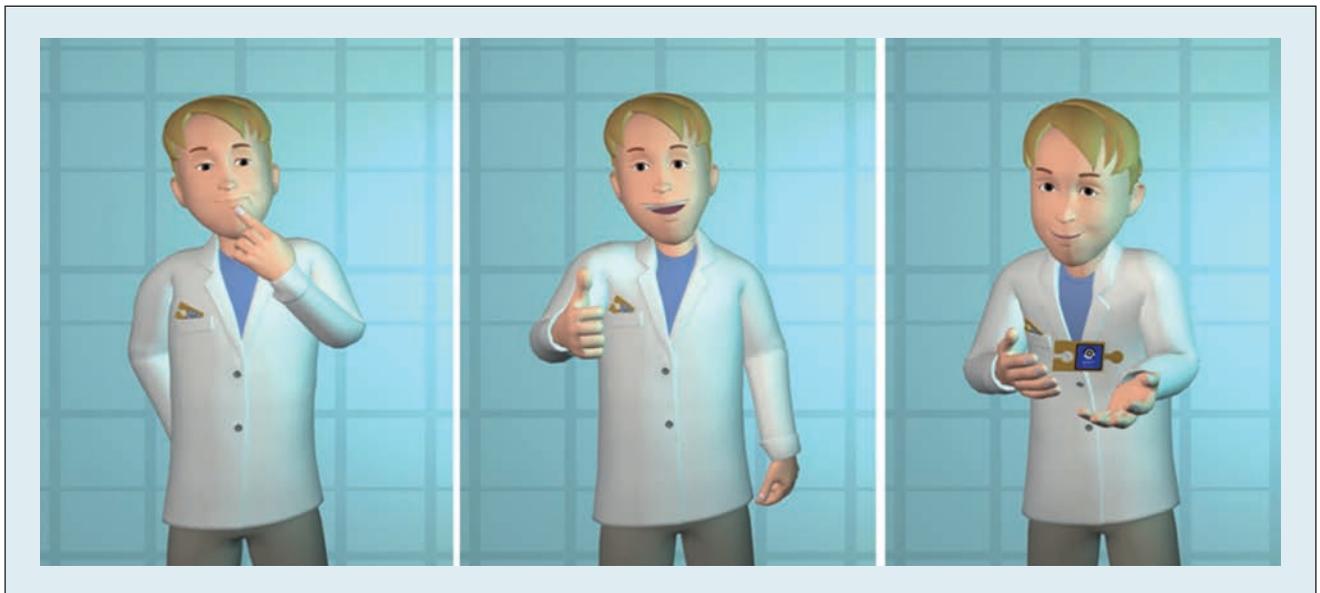


Figure 1. Coach Mike, a pedagogical agent, in three different poses.<sup>7</sup> Source: H. Chad Lane; used with permission.

robotics. This kind of application extends the role of classroom teaching<sup>7</sup>:

It means that such systems need to go beyond simply focusing on knowledge outcomes. They must take seriously goals such as convincing a visitor to engage, promoting curiosity and interest, and ensuring that a visitor has a positive learning experience. In other words, pedagogical agents for informal learning need to not only act as coach (or teacher), but also as advocate (or salesperson).

Coach Mike was designed to emulate some of the human museum cura-

tors' work, including helping orientate visitors, encouraging them to explore, and providing problem-solving challenges and support.

Some researchers have recently argued the benefits of AI systems in education,<sup>8,9</sup> whereas others have been more skeptical.<sup>10</sup> This column looks at the evidence derived from meta-reviews and meta-analyses conducted over the past five years. Its main focus is on the comparative effectiveness of AIED systems versus human tutoring. Note that a meta-review of the use of pedagogical agents (not necessarily in AIED systems) “pro-

duced a small but significant effect on learning.”<sup>11</sup>

This column is not intended as support for an argument about getting rid of human teachers, but rather as support for blended learning, in which some of the human teacher's work can be offloaded to AIED systems, as if to a classroom assistant.

### Meta-Analysis and Metareviews

Since 2011, several meta-reviews and meta-analyses have attempted to determine the degree to which a whole host of systems have been educationally

**Table 1. Effect sizes adapted from work by Kurt VanLehn.<sup>12</sup>**

Comparison	No. studies	Mean effect size	Reliability (%)
Answer based vs. no tutoring*	165	0.31	40
Step based vs. no tutoring	28	0.76	68
Substep based vs. no tutoring	26	0.40	54
Human vs. no tutoring	10	0.79	80
Step based vs. answer based	2	0.40	50
Substep based vs. answer based	6	0.32	33
Human vs. answer based	1	-0.04	0
Substep based vs. step based	11	0.16	0
Human vs. step based	10	0.21	30
Human vs. substep based	5	-0.12	0

\*Row 1 was taken by VanLehn from a separate study.<sup>13</sup>

effective. Typically, this has meant comparing them in terms of learning gains with other instructional methods, such as whole-class teaching by a human teacher or the use of a textbook without a teacher.

### VanLehn's Meta-Analysis

Kurt VanLehn analyzed papers comparing five types of tutoring: no tutoring (for example, learning with just a textbook), answer-based tutoring, step-based tutoring, substep-based tutoring, and human tutoring.<sup>12</sup>

The difference between answer-based, step-based, and substep-based tutoring resides in the granularity of the interaction between the tutor and student (see the sidebar, “Answer-, Step-, and Substep-Based Tutoring”). Answer-based systems can provide hints and feedback only at the level of the overall answer. Step-based systems can provide hints, scaffolding, and feedback on every step that the student makes in solving the problem. By contrast, substep-based systems work at a finer granularity level still and “can give scaffolding and feedback at a level of detail that is even finer than the steps students would normally enter when solving a problem.”<sup>12</sup> AI techniques are required to underpin both step-based and substep-based tutors, whereas answer-based systems would typically fall under the heading of computer-based or computer-assisted instruction (CAI).

Given these granularity levels, VanLehn derived 10 pairwise comparisons of effect sizes (see Table 1). The rightmost column shows the proportion of the results for that row where the individual study comparison was statistically reliable at the level  $p < 0.05$ .

For the purposes of this review, the most interesting comparison is that between one-on-one human tutoring and step-based tutors (effect size = 0.21). By collating all the results in Table 1, VanLehn found that step-based tutors were, within the limitations of his review, “just as effective as adult, one-on-one tutoring for increasing learning gains in STEM topics.”<sup>12</sup> He also found that although increasing the granularity of instruction from answer-based to step-based yielded significant gains, going to the finer level of substep-based tutoring did not add further value. (Note that this latter finding was based on a small number of studies only.)

### Six Metareviews

Since VanLehn's meta-analysis, six metareviews have been published, as well as a large-scale study of a specific tutor (see Table 2). In the table, the “number of comparisons” column shows the number of instances for the given comparison in that row, not the total number of studies in the overall metareview.

In a metareview of 107 studies, Wenting Ma and colleagues found similar results to VanLehn for step-based ITSs both when compared to a no-tutoring condition (that is, just a textbook; mean effect size = 0.36) and, more positively than VanLehn, when compared to large-group instruction led by a human teacher (mean effect size = 0.44).<sup>14</sup> They found no differences when compared to small group human tutoring or one-on-one tutoring.

The same authors analyzed 22 systems for teaching programming and also found a “significant advantage of ITS over teacher-led classroom instruction and non-ITS computer-based instruction.”<sup>15</sup> A larger version of a similar study involving 280 studies is currently in progress.<sup>16</sup>

In a metareview of 50 studies involving 63 comparisons, James Kulik and J.D. Fletcher found comparable improvements (mean effect size = 0.65),<sup>17</sup> but they distinguished studies that used standardized tests from those in which the tests were more specifically tuned to the system providing tuition, with smaller effect sizes when standardized tests were employed. Overall, they concluded that “this meta-analysis shows that ITSs can be very effective instructional tools ... Developers of ITSs long ago set out to improve on the success of CAI tutoring and to match the success of human tutoring. Our results suggest that ITS developers have already met both of these goals.”<sup>17</sup> They also found better results for substep-based systems than VanLehn, which they ascribed to differing comparison methodologies.

Much smaller effect sizes were found by Saiying Steenbergen-Hu and Harris Cooper in their meta-analysis of pupils using ITSs in school settings.<sup>18</sup> Kulik and Fletcher put this down to the weaker study inclusion criteria (for example, the inclusion of answer-based systems as if they were

**Table 2. Six metareviews and a large-scale study.**

Row number	Metareview	Comparison	No. comparisons	Mean effect size	Standard error
1	VanLehn <sup>12</sup>	Step based vs. one-on-one human tutoring	10	-0.21	0.19 <sup>†</sup>
2	Wenting Ma and colleagues <sup>14</sup>	Step based vs. one-on-one human tutoring	5	-0.11	0.10
3	Ma and colleagues <sup>14</sup>	Step based vs. large group human instruction	66	0.44	0.05
4	John Nesbit and colleagues <sup>15</sup>	Step based vs. teacher-led group instruction	11	0.67	0.09
5	James Kulik and J.D. Fletcher <sup>17</sup>	Step based and substep based vs. conventional classes	63	0.65	0.07 <sup>†</sup>
6	Saiying Steenbergen-Hu and Harris Cooper (2014) <sup>19</sup>	Step based vs. one-on-one human tutoring	3	-0.25	0.24
7	Steenbergen-Hu and Cooper (2014) <sup>19</sup>	Step based vs. traditional classroom instruction	16	0.37	0.07
8	Steenbergen-Hu and Cooper (2013) <sup>18</sup>	Step based and answer based vs. traditional classroom instruction	26	0.09	0.01
9	John Pane and colleagues <sup>20</sup>	Blended learning including a step-based system vs. traditional classroom instruction	147 schools	-0.1	0.10
				<b>0.21</b>	<b>0.10</b>
				0.01	0.11
				0.19	0.14
10	Weighted mean	AIED system vs. one-on-one human tutoring	18	-0.19	N/A
11	Weighted mean	AIED system vs. conventional classes	182	0.47	N/A

\* The standard error in row 1 is based on all 10 studies, not just the 30% that produced reliable results (see Table 1).

† Standard errors computed by this article's author.

step-based systems) used by Steenbergen-Hu and Cooper, who also noted that lower achievers seemed to do worse with ITSs than did the broad spectrum of school pupils, although Kulik and Fletcher disputed this result.<sup>17</sup> However, in a parallel study of university students, Steenbergen-Hu and Cooper found more positive effect sizes (in the range of 0.32 to 0.37) for ITSs as compared to conventional instruction.<sup>19</sup> They conclude that ITSs “have demonstrated their ability to outperform many instructional methods or learning activities in facilitating college-level students’ learning of a wide range of subjects, although they are not as effective as human tutors. ITSs appear to have a more pronounced effect on college-level learners than on K–12 students.”<sup>19</sup>

Rows 10 and 11 of Table 2 summarize the results of the metareviews, excluding the evaluation of the Cognitive Algebra Tutor (row 9), and show a weighted mean effect size of 0.47 for AIED systems versus conventional classroom teaching. We use the term *AIED system* to cover all the systems—step-based, substep-based, and answer-based—looked at in the metareviews. The comparison with one-on-one human tutoring shows that AIED systems do slightly worse, with a mean effect size of -0.19. In both cases, the means are weighted in terms of the number of comparisons in the metareview, not in terms of the original *N* values in the studies themselves.

### Cognitive Tutors

Cognitive Tutors “are found in about 3,000 schools, and over a half mil-

lion students use the courses each year.”<sup>21</sup> They represent the most successful transition, in terms of numbers of students, of AIED work from the laboratory to the classroom. They provide scaffolded help with step-by-step problem-solving in various domains, mostly mathematical, and are designed to be used in a blended learning manner, thus freeing up the teacher to work with other children. Teachers are trained to make the best of these systems’ arrival in their classrooms in terms of managing all the pupils in the classroom before, during, and after the use of the tutors.<sup>4</sup> Individual evaluations of various Cognitive Tutors are included in the reviews already described.

A large-scale US study of the Cognitive Algebra Tutor undertook a

between-schools project involving 73 high schools and 74 middle schools across seven states.<sup>20</sup> The schools were matched in pairs: half received the Cognitive Algebra Tutor and adjusted their teaching to include it as they saw fit, whereas the others carried on with their regular method of teaching algebra. The study ran for two years and found no significant differences on post-test scores in the first year of the study, but in the second year, the high schools that used the Cognitive Tutor showed a small but significant effect size of 0.21 (see the bolded data in row 9 of Table 2).

Note that how the Cognitive Tutor was actually used in the classrooms was not controlled, although post hoc analyses showed that teachers did not generally use the Tutor exactly as recommended by its developers.

**T**he overall conclusion of these meta-reviews and analyses is that AIED systems perform better than both CAI systems and human teachers working in large classes. They perform slightly worse than one-on-one human tutors. Most of the systems taught mathematics or STEM subjects, because these are the kinds of subject for which it is easier to build the domain and student models mentioned in the introduction. Note that there was a degree of overlap between these meta-reviews and analyses in terms of the collections of individual evaluations from which they have drawn their conclusions.

The specific study of the Cognitive Tutor for Algebra evaluated its use as a blended addition to the regular algebra teaching in the schools in which it was tried, rather than as a total replacement for the teachers, and found good results in high schools as opposed to mid-

dle schools and in the second year of the evaluation as opposed to the first year. For various reasons, the way forward for AIED systems in the classroom must be the blended model—classroom assistants, if you like—in order to provide detailed one-on-one tutoring for some students while the human teacher attends to others, as well as having overall responsibility for all the students' progress.

Of course, good post-test results are not the only criteria for judging whether an educational technology will be or should be adopted.<sup>10</sup> However, the overall message of these evaluations is that blending AIED technology with other forms of teaching is beneficial, particularly for older pupils and college-level students studying STEM subjects. ■

### Acknowledgments

This column is an adapted and expanded version of a letter to the editor of the *International Journal of Artificial Intelligence in Education*.<sup>22</sup>

### References

1. J.S. Brown, R.R. Burton, and A.G. Bell, "SOPHIE: A Step Towards a Reactive Learning Environment," *Int'l J. Man-Machine Studies*, vol. 7, 1975, pp. 675–696.
2. T. O'Shea, "A Self-Improving Quadratic Tutor," *Intelligent Tutoring Systems*, D. Sleeman and J.S. Brown, eds., Academic Press, 1982.
3. J.S. Brown and R.R. Burton, "Multiple Representations of Knowledge for Tutorial Reasoning," *Representation and Understanding*, D.G. Bobrow and A. Collins, eds., Academic Press, 1975, pp. 311–349.
4. K.R. Koedinger et al., "Intelligent Tutoring Goes to School in the Big City," *Int'l J. Artificial Intelligence in Education*, vol. 8, no. 1, 1997, pp. 30–43.

5. K. Leelawong and G. Biswas, "Designing Learning by Teaching Agents: The Betty's Brain System," *Int'l J. Artificial Intelligence in Education*, vol. 18, no. 3, 2008, pp. 181–208.
6. W.L. Johnson, "Serious Use of a Serious Game for Language Learning," *Int'l J. Artificial Intelligence in Education*, vol. 20, no. 2, 2010, pp. 175–195.
7. H.C. Lane et al., "The Effects of a Pedagogical Agent for Informal Science Education on Learner Behaviors and Self-Efficacy," *Proc. 16th Int'l Conf. Artificial Intelligence in Education*, 2013, pp. 309–318.
8. R. Luckin et al., *Intelligence Unleashed: An Argument for AI in Education*, Pearson, 2016.
9. B.P. Woolf et al., "AI Grand Challenges for Education," *AI Magazine*, vol. 34, no. 4, 2014, pp. 66–84.
10. N. Enyedy, "Personalized Instruction: New Interest, Old Rhetoric, Limited Results, and the Need for a New Direction for Computer-Mediated Learning," policy brief, National Education Policy Center, 2014.
11. N.L. Schroeder, O.O. Adesope, and R.B. Gilbert, "How Effective Are Pedagogical Agents for Learning? A Meta-Analytic Review," *J. Educational Computing Research*, vol. 49, no. 1, 2013, pp. 1–39.
12. K. VanLehn, "The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems," *Educational Psychologist*, vol. 46, no. 4, 2011, pp. 197–221.
13. C.-L.C. Kulik and J.A. Kulik, "Effectiveness of Computer-Based Instruction: An Updated Analysis," *Computers in Human Behavior*, vol. 7, nos. 1–2, 1991, pp. 75–94.
14. W. Ma et al., "Intelligent Tutoring Systems and Learning Outcomes: A Meta-Analysis," *J. Educational Psychology*, vol. 106, no. 4, 2014, pp. 901–918.
15. J.C. Nesbit et al., "How Effective Are Intelligent Tutoring Systems in Computer Science Education?" *Proc. IEEE 14th Int'l Conf. Advanced Learning Technologies*, 2014, pp. 99–103.

16. J.C. Nesbit et al., "Work in Progress: Intelligent Tutoring Systems in Computer Science and Software Engineering Education," *Proc. 122nd Am. Soc. Eng. Education Ann. Conf.*, 2015, paper 11861.
17. J.A. Kulik and J.D. Fletcher, "Effectiveness of Intelligent Tutoring Systems: A Meta-Analytic Review," *Rev. Educational Research*, vol. 86, no. 1, 2016, pp. 42–78.
18. S. Steenbergen-Hu and H. Cooper, "A Meta-Analysis of the Effectiveness of Intelligent Tutoring Systems on K–12 Students' Mathematical Learning," *J. Educational Psychology*, vol. 105, no. 4, 2013, pp. 970–987.
19. S. Steenbergen-Hu and H. Cooper, "A Meta-Analysis of the Effectiveness of Intelligent Tutoring Systems on College Students' Academic Learning," *J. Educational Psychology*, vol. 106, no. 2, 2014, pp. 331–347.S.
20. J.F. Pane et al., "Effectiveness of Cognitive Tutor Algebra I at Scale," *Educational Evaluation and Policy Analysis*, vol. 36, no. 2, 2014, pp. 127–144.
21. K.R. Koedinger and V. Alevan, "An Interview Reflection on 'Intelligent Tutoring Goes to School in the Big City,'" *Int'l J. Artificial Intelligence in Education*, vol. 16, no. 1, 2016, pp. 13–24.
22. B. du Boulay, "Recent Meta-Reviews and Meta-Analyses of AIED Systems," *Int'l J. Artificial Intelligence in Education*, vol. 26, no. 1, 2016, pp. 536–537.

**Benedict du Boulay** is an emeritus professor of artificial intelligence and a member of the Human Centered Technology Research Group at the University of Sussex. Contact him at [b.du-boulay@sussex.ac.uk](mailto:b.du-boulay@sussex.ac.uk).

*This article originally appeared in IEEE Intelligent Systems, vol. 31, no. 6, 2016.*

## IEEE computer society

**PURPOSE:** The IEEE Computer Society is the world's largest association of computing professionals and is the leading provider of technical information in the field.

**MEMBERSHIP:** Members receive the monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.

**OMBUDSMAN:** Email [ombudsman@computer.org](mailto:ombudsman@computer.org).

**COMPUTER SOCIETY WEBSITE:** [www.computer.org](http://www.computer.org)

**Next Board Meeting:** 12–17 June 2017, Phoenix, AZ, USA

### EXECUTIVE COMMITTEE

**President:** Jean-Luc Gaudiot; **President-Elect:** Hironori Kasahara; **Past President:** Roger U. Fujii; **Secretary:** Forrest Shull; **First VP, Treasurer:** David Lomet; **Second VP, Publications:** Gregory T. Byrd; **VP, Member & Geographic Activities:** Cecilia Metra; **VP, Professional & Educational Activities:** Andy T. Chen; **VP, Standards Activities:** Jon Rosdahl; **VP, Technical & Conference Activities:** Hausi A. Müller; **2017–2018 IEEE Director & Delegate Division VIII:** Dejan S. Milošević; **2016–2017 IEEE Director & Delegate Division V:** Harold Javid; **2017 IEEE Director-Elect & Delegate Division V-Elect:** John W. Walz

### BOARD OF GOVERNORS

**Term Expiring 2017:** Alfredo Benso, Sy-Yen Kuo, Ming C. Lin, Fabrizio Lombardi, Hausi A. Müller, Dimitrios Serpanos, Forrest J. Shull  
**Term Expiring 2018:** Ann DeMarle, Fred Douglass, Vladimir Getov, Bruce M. McMillin, Cecilia Metra, Kunio Uchiyama, Stefano Zanero  
**Term Expiring 2019:** Saurabh Bagchi, Leila De Floriani, David S. Ebert, Jill I. Gostin, William Gropp, Sumi Helal, Avi Mendelson

### EXECUTIVE STAFF

**Executive Director:** Angela R. Burgess; **Director, Governance & Associate Executive Director:** Anne Marie Kelly; **Director, Finance & Accounting:** Sunny Hwang; **Director, Information Technology & Services:** Sumit Kacker; **Director, Membership Development:** Eric Berkowitz; **Director, Products & Services:** Evan M. Butterfield; **Director, Sales & Marketing:** Chris Jensen

### COMPUTER SOCIETY OFFICES

**Washington, D.C.:** 2001 L St., Ste. 700, Washington, D.C. 20036-4928  
**Phone:** +1 202 371 0101 • **Fax:** +1 202 728 9614 • **Email:** [hq.ofc@computer.org](mailto:hq.ofc@computer.org)  
**Los Alamitos:** 10662 Los Vaqueros Circle, Los Alamitos, CA 90720  
**Phone:** +1 714 821 8380 • **Email:** [help@computer.org](mailto:help@computer.org)

### MEMBERSHIP & PUBLICATION ORDERS

**Phone:** +1 800 272 6657 • **Fax:** +1 714 821 4641 • **Email:** [help@computer.org](mailto:help@computer.org)  
**Asia/Pacific:** Watanabe Building, 1-4-2 Minami-Aoyama, Minato-ku, Tokyo 107-0062, Japan • **Phone:** +81 3 3408 3118 • **Fax:** +81 3 3408 3553 • **Email:** [tokyo.ofc@computer.org](mailto:tokyo.ofc@computer.org)

### IEEE BOARD OF DIRECTORS

**President & CEO:** Karen Bartleson; **President-Elect:** James Jefferies; **Past President:** Barry L. Shoop; **Secretary:** William Walsh; **Treasurer:** John W. Walz; **Director & President, IEEE-USA:** Karen Pedersen; **Director & President, Standards Association:** Forrest Don Wright; **Director & VP, Educational Activities:** S.K. Ramesh; **Director & VP, Membership and Geographic Activities:** Mary Ellen Randall; **Director & VP, Publication Services and Products:** Samir El-Ghazaly; **Director & VP, Technical Activities:** Marina Ruggieri; **Director & Delegate Division V:** Harold Javid; **Director & Delegate Division VIII:** Dejan S. Milošević

revised 26 Jan. 2017





# From the Editor in Chief

Editor in Chief: Maria R. Ebling ■ IBM T.J. Watson Research Center ■ ebling@us.ibm.com

## TJBot and Zooids: The Connection between Pervasive Computing and AI

Maria R. Ebling, IBM T.J. Watson Research Center

**P**ervasive computing and AI have always struck me as a match made in heaven. The notion that technology will fade into the background of our lives necessarily requires understanding and anticipating our needs, which requires some level of AI.

### BRINGING AI TO PERVASIVE COMPUTING

The maker community has recently rolled out kits to bring these two areas of computer science together. Two notable examples are the TJBot and the Zooids.

TJBot is a DIY kit for creating a cardboard robot with Watson inside. Inside TJBot is a Raspberry Pi with a camera, microphone, speaker, and LED light with easy connection to Watson services. The TJBot kit provides three starter recipes. The first is a sentiment analysis that changes the color of TJBot's LED based on the emotional sentiment of a particular topic on Twitter. The second lets you use your voice to control TJBot. The third lets you have a conversation with TJBot. This open source project is available on GitHub. Directions for printing your very own TJBot, as well as the initial recipes and recipes contributed by the community, are available at <https://ibmtjbot.github.io>.

### EDITORIAL BOARD CHANGES

I'm sorry to announce that Anind Dey will be stepping down as Associate Editor-in-Chief. His expertise and leadership in the area of human-computer interaction, especially in context-aware computing, will be missed. We thank him for his many years of service to *IEEE Pervasive Computing*.

Zooids are small, ant-like robots with little wheels, a touch sensor, gyros, and an optical sensor on top to monitor instructions from a projector overhead. Zooids work cooperatively, in a swarm, to accomplish things like moving your phone closer to you or coming out of a mouse hole in the wall at night to tidy up while you sleep (I like this one the best but it's still just a vision). Like TJBot, Zooids software is available on GitHub, and the authors invite interested community members to try them out. You can find the software at <https://github.com/ShapeLab/SwarmUI>.

Interestingly, both of these maker projects are exploring fundamental questions about human-computer interaction. TJBot is looking to understand how humans will (want to) interact with cognitive objects, especially using voice commands. Zooids are exploring how humans might exploit swarm-based user interfaces consisting of small, tangible robots working in concert with one another and a human.

Both of these projects also leverage internal AI capabilities—either explicitly or implicitly. These kits highlight that our industry is on the cusp of fundamental changes in the way we will interact with technology.

### IN THIS ISSUE

Drones, our theme for this issue, are another good example of technology that falls at the intersection of pervasive computing with AI. Many thanks to Albrecht Schmidt and Floyd Mueller for focusing our attention on drones. In this issue, we also introduce two new departments: IoT News and Human Augmentation.

Our first new department, IoT News, is led by Florian Michahelles from Siemens and presents news and research on the Internet of Things. The coverage will range from industrial use cases of IoT to social and business implications. Michahelles is looking for contributions from the community and outlines, in more detail, what he is looking for

**MISSION STATEMENT:** *IEEE Pervasive Computing* is a catalyst for advancing research and practice in mobile and ubiquitous computing. It is the premier publishing forum for peer-reviewed articles, industry news, surveys, and tutorials for a broad, multidisciplinary community.



*IEEE Intelligent Systems* delivers the latest peer-reviewed research on all aspects of artificial intelligence, focusing on practical, fielded applications. Contributors include leading experts in

- Intelligent Agents • The Semantic Web
- Natural Language Processing
- Robotics • Machine Learning

Visit us on the Web at  
[www.computer.org/intelligent](http://www.computer.org/intelligent)



## Call for Articles

*IEEE Software* seeks practical, readable articles that will appeal to experts and nonexperts alike. The magazine aims to deliver reliable, useful, leading-edge information to software developers, engineers, and managers to help them stay on top of rapid technology change. Topics include requirements, design, construction, tools, project management, process improvement, maintenance, testing, education and training, quality, standards, and more.

Author guidelines:

[www.computer.org/software/author](http://www.computer.org/software/author)

Further details: [software@computer.org](mailto:software@computer.org)

[www.computer.org/software](http://www.computer.org/software)



from community submissions. Please read through it and consider submitting material for future issues.

Our second new department, Human Augmentation, is led by Albrecht Schmidt. In this department, Schmidt plans to discuss technologies designed to augment the human intellect, including cognition and perception. In this issue, he argues that the ongoing technical revolution is the modern equivalent of the industrial revolution. He explains how today's technologies enhance our cognitive abilities much the way technologies of the industrial revolution enhanced our physical abilities. This department will help us understand the societal implications of pervasive computing and IoT.

In our Smartphones department, Rajalakshmi Nandakumar and Shyam-nath Gollakota discuss the prototype system they built, demonstrating some compelling use cases for active sonar implementations on smartphones. They describe two major use cases. First, they built and studied a system that detects sleep disorders using active sonar on a smartphone. What makes this compelling is that you no longer have to attach all kinds of wires to the patient using white goo in their hair! Second, they

built a system that follows human gestures from the phone, even when the phone is in your pocket. Imagine being able to use gestures to control your phone without having to dig it out of your pocket! I found the article fascinating and seemingly right out of a science fiction novel.

In our Pervasive Health department, Kay Connelly, Oscar Mayora, Jesus Favela, Maia Jacobs, Aleksandar Matic, Chris Nugent, and Stefan Wagner report on the outcomes of a set of discussions held at the Future of Pervasive Health Workshop. During this workshop, they identified five major themes: technological challenges and opportunities, adoption and adherence of pervasive health solutions, open data for pervasive health, pervasive health methods and ethical issues, and the road to education on pervasive health. I found the summary of the discussion very interesting, especially the need for open data and a formal curriculum.

Finally, our Notes from the Community department highlights some really interesting technologies. As someone who has experienced physical therapy, including for repetitive strain injuries of the arms, I found the skintillates temporary

## FROM THE EDITOR IN CHIEF

tattoos particularly intriguing. The potential of this technology for use in physical therapy seems enormous! My physical therapist could see the potential in helping change people's habits, especially when they sit in front of a keyboard.

We also have three feature articles this issue, starting with "A Survey of Diet Monitoring Technology," by Haik Kalantarian, Nabil Alshurafa, and Majid Sarrafzadeh. This article provides a deep-dive analysis of different technologies available for monitoring eating behaviors. These technologies range from manual tracking, to acoustic monitoring, to gesture recognition, to instrumented objects, to visual approaches, to sensors attached to the skin. The authors also performed a survey designed to measure user acceptance of these technologies. This article provides a thorough review of the many ways to do automated tracking of eating behaviors.

The next article is "Semi-Automated Tracking: A Balanced Approach for Self-Monitoring Applications," by Eun Kyoung Choe, Saeed Abdullah, Mashfiqui Rabbi, Edison Thomaz, Daniel A. Epstein, Felicia Cordeiro, Matthew Kay, Gregory D. Abowd, Tanzeem Choudhury, James Fogarty, Bongshin Lee, Mark Matthews, and Julie A. Kientz. The authors make the case for semi-automated tracking of health behaviors. Semi-automated tracking balances the convenience of automated tracking with the awareness of manual tracking. The authors define semi-automated tracking and walk readers through its characteristics, contrasting the strengths and weaknesses with those of both manual and automated tracking. They then examine three important design considerations, and give examples of semi-automated tracking of food, mood, and sleep to demonstrate these ideas. Semi-automated tracking strikes me as a promising compromise between fully automated and fully manual tracking approaches.

Finally, in "Region Formation for Efficient Offline Location Prediction," Ian Craig and Mark Whitty explore the use of regions to make predictions

about users' destinations. The authors recognize the importance of location predictions as well as the challenges faced in making such predictions on mobile devices, including power consumption, privacy, and data sparsity. To overcome these challenges, they propose dividing the geographic area into regions that uniquely differentiate users' paths. This approach addresses the power consumption problem because updates are needed much less frequently. It addresses the privacy problem because the user's exact location within the region is unimportant. It addresses the data sparsity problem by reducing the overall state space. Although this approach shows promise, more research is needed before it is applied in "the real world."

The ability to bring AI together with pervasive computing will help us to create more compelling user experiences in the years to come. As we move forward, I anticipate these two fields coming together more and more, in applications from health, to drones, to smarter cities. ■

**Maria R. Ebling** is a director at the IBM T.J. Watson Research Center. She manages a team building systems capable of supporting a Smarter Planet while not forgetting about the people who use such systems. Ebling received her PhD in computer science from Carnegie Mellon University. She's a member of the IBM Academy of Technology, a distinguished member of the ACM, and a senior member of IEEE. Contact her at [ebbling@us.ibm.com](mailto:ebbling@us.ibm.com).

*This article originally appeared in  
IEEE Pervasive Computing, vol. 16, no. 1, 2017.*

**Call for Articles**

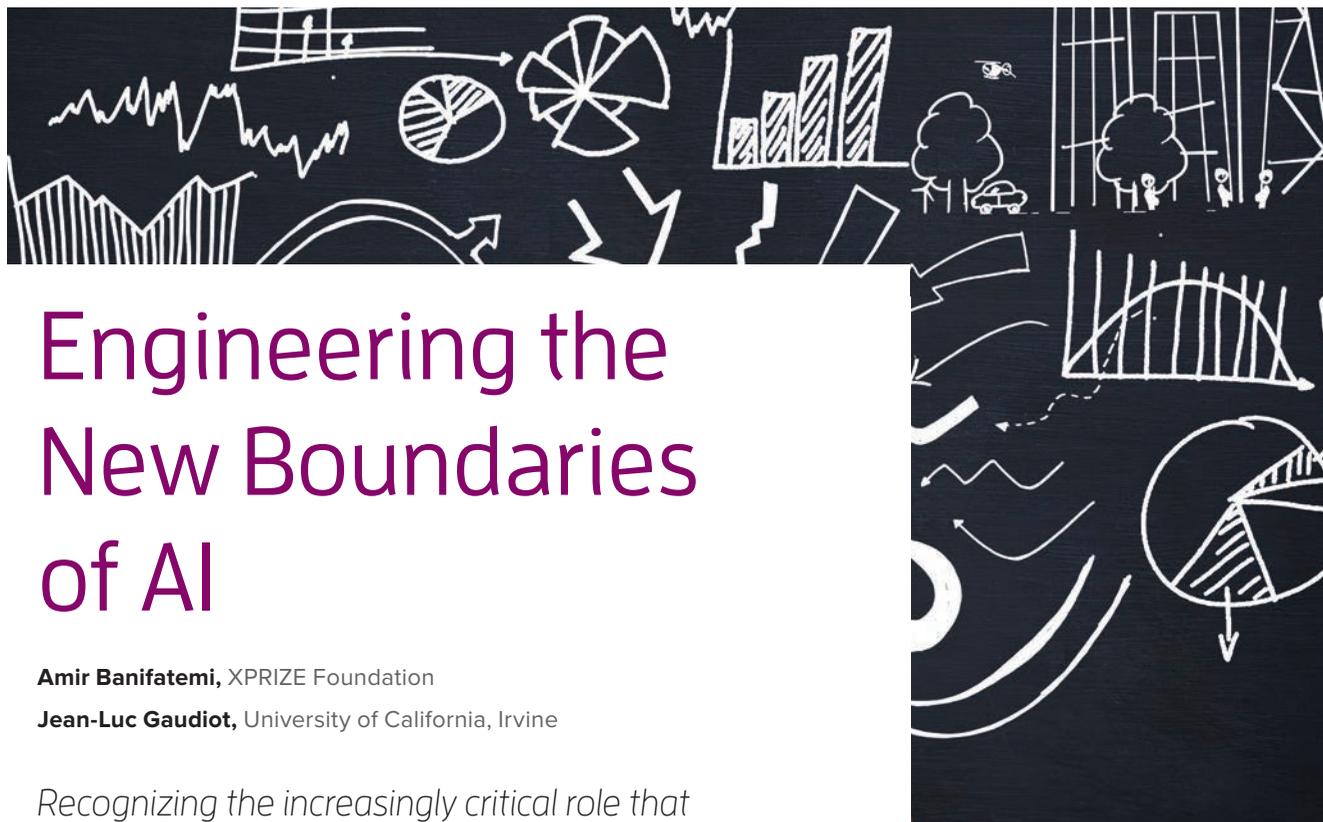
**IEEE Pervasive Computing**

seeks accessible, useful papers on the latest peer-reviewed developments in pervasive, mobile, and ubiquitous computing. Topics include hardware technology, software infrastructure, real-world sensing and interaction, human-computer interaction, and systems considerations, including deployment, scalability, security, and privacy.

**Further details:**  
[pervasive@computer.org](mailto:pervasive@computer.org)  
[www.computer.org/pervasive](http://www.computer.org/pervasive)

**Author guidelines:**  
[www.computer.org/mc/pervasive/author.htm](http://www.computer.org/mc/pervasive/author.htm)

**IEEE pervasive COMPUTING**  
 MOBILE AND UBIQUITOUS SYSTEMS



# Engineering the New Boundaries of AI

**Amir Banifatemi**, XPRIZE Foundation

**Jean-Luc Gaudiot**, University of California, Irvine

*Recognizing the increasingly critical role that AI plays in all aspects of modern society, the XPRIZE Foundation launched the IBM Watson AI XPRIZE. Interdisciplinary teams will advance current AI technologies as they compete for the grand prize.*

**A**I surrounds us. It's in our search engines, automobiles, phones, video-streaming websites, and even the systems that run our financial markets. We interact with and rely on intelligent machines throughout our daily activities.

AI has made steady, linear progress toward adoption and integration over the past 50 years. It's now at an important inflection point—the adoption curve could soon transcend the linear growth pattern and take an exponential leap. With machine learning, AI technologies have the potential to make sense of vast, complex datasets. This begs the question: How can we best leverage both human intelligence and these powerful AI technologies to benefit humanity now and in the future?

\$3 million, a second-place prize of \$1 million, and a third-place prize of \$500,000. The goal of the competition is to accelerate the development of scalable AI solutions to address humanity's greatest challenges. The four-year competition will have milestone completions in 2017, 2018, and 2019, with the top three finalists competing for the grand prize in 2020.

Recognizing the competition's humanitarian importance and global reach, the IEEE Computer Society (CS) formed a partnership agreement in August 2016 with the XPRIZE Foundation. Select CS members will join the XPRIZE scientific advisory board, and the CS will draw from its global network of experts to help judge and advise the competitors. The CS will also encourage

## ENTER XPRIZE

It's this question of the impact of AI on mankind that inspired the IBM Watson AI XPRIZE, the latest innovation challenge from the XPRIZE Foundation. IBM sponsored the competition with a prize purse of \$5 million, including a grand prize of

## XPRIZE ORIGIN STORY

The history of XPRIZE is unique. The first was inspired by the \$25,000 Orteig Prize, which was awarded in 1927 to American aviator Charles Lindbergh for flying the *Spirit of St. Louis* on the first nonstop flight between New York and Paris.

The original XPRIZE, announced in 1996, was a \$10 million prize to the first privately financed team that could build and fly a three-passenger vehicle 100 km into space twice within two weeks. The prize, later titled the Ansari XPRIZE for suborbital spaceflight, motivated 26 teams from 7 nations to invest more than \$100 million in pursuit of the \$10 million prize. On 4 October 2004, the prize was awarded to Mojave Aerospace Ventures, kicking off the personal spaceflight revolution.

The first competition created a funding model to stimulate broad investment in innovation that produces 10-fold return on the prize as well as a 100-fold increase in follow-on investment and social benefit.

The XPRIZE Foundation recognized that this extraordinary leverage could be beneficial for a range of global grand challenges in which these inducement prizes would lead to drastic and important improvements in quality of life.

Active competitions include the \$30 million Google Lunar XPRIZE, the \$20 million NRG COSIA Carbon XPRIZE, the \$15 million Global Learning XPRIZE, the \$10 million Qualcomm Tricorder XPRIZE, the \$7 million Shell Ocean Discovery XPRIZE, the \$7 million Barbara Bush Foundation Adult Literacy XPRIZE, and the \$5 million IBM Watson AI XPRIZE.

its members to serve as mentors or to provide other technical expertise to entrants in support of improving their application or approach as well as in showcasing their team's progressive achievements.

The XPRIZE Foundation launched its first competition in 1996 in an effort to spur privately funded innovation and research for space exploration (learn more about XPRIZE's history in the "XPRIZE Origin Story" sidebar). With the success of this first competition, the foundation has attracted continued private funding for additional XPRIZE competitions in health, learning, energy, environment, global development and exploration.

These privately funded, prize-based scientific challenges are highly attractive to individuals, companies, and organizations, and they're clearly inspiring a solution-focused approach to grand challenges that is less research oriented than traditional

academic or corporate settings. University- and corporation-based research and development traditionally have incentivized research advances through promotion- or salary-based structures. XPRIZE funding and publicity are untethered to any one corporation or institution, so they offer global recognition and opportunities for successful competitors.

XPRIZE competitions draw significant numbers of nontraditional innovators and entrepreneurs and enable them to focus completely on one area of innovation. With its goal to bring about "radical breakthroughs for the benefit of humanity" through incentivized competition, the XPRIZE approach doesn't impose budget requirements or other significant administrative overhead. Researchers, engineers, and entrepreneurs are free from constraints that might typically limit them with government- or consortia-based funding sources. By

driving competitors to invest every intellectual and financial resource at their disposal to solve the problem, reach the goal, and win the prize, innovation can take center stage. A large number of groups can work in parallel toward the goal, and the resulting solution or benefit can be quickly adopted and realized by society.

## INCENTIVE-DRIVEN COMPETITIONS SPUR INNOVATION

Challenge-based models offering substantial financial prizes have historically succeed in driving innovation. This departure from traditional research models enables competitors to focus on a solution's impact, rather than the particulars of the methods that are used to achieve it. Contributors are inspired to explore all possible and creative ways to meet that objectives and win.

Of course, prizes for innovation and technological achievements aren't new: The British government's simple and practical method to precisely determine a ship's longitude; Charles Lindbergh won the Orteig Prize for the first transatlantic flight from New York to Paris; and, more recently, Elbert "Burt" Rutan's team won the Ansari XPRIZE for developing the first reusable private spaceship.

With tech-focused prizes, awards are tied to achievements versus attempts, and teams are driven to leverage any and all resources and support, pulling together an impressive matrix of expertise, capital, collaboration support, and achievement-oriented structure. In addition, recent incentive-driven competitions have shown that teams invest their own resources in these competitions to the tune of 10 to 100 times the prize itself.

Competitions can come in all sizes—from the smaller Kaggle Competitions ([kaggle.com/competition](http://kaggle.com/competition)) to the larger \$1 million Netflix Prize—thus there's more awareness and acceptance of such incentive prizes than ever. Of course, scale and size matter, and

multimillion-dollar competitions tend to have a greater impact on the market and innovation than smaller, more localized versions. However, their greater popularity has brought them increasingly into other domains—including the classroom. Educators have been enthusiastic to explore challenge-based learning, appreciating their potential for developing students' abilities to work in teams, to be inventive and resourceful, and to push beyond traditional lines of thought and institutional boundaries. Indeed, leveraging such challenges in the learning environment imparts some unique skillsets that are invaluable for students as they transition to the modern workplace.

## THE AI XPRIZE

As the first “open” challenge in AI, the IBM Watson AI XPRIZE is unlike its predecessors. Rather than set a single universal goal for all teams, this competition invites teams to each create their own goal: an application of AI to a meaningful problem. Once goals are defined, the teams will use their own best available resources, along with assistance from selected experts in the professional AI and engineering communities, to develop their plans for the duration of the competition.

Each team must demonstrate progress and annual milestone achievements for their plan, which will ultimately qualify them to compete for the grand prize on the TED 2020 stage in front of a live in-person and online audience. The audience and a panel of expert judges will crown the grand-prize winner and rank the runners-up from the three finalists.

As AI encompasses a wide range of disciplines, teams can utilize several types of specialized technology in achieving their goals. Technologies that drive or support intelligent behaviors also integrate sensory inputs and cognitive capabilities. These systems support human-machine collaboration through their ability to possess domain understanding, learn from

## JOIN THE COMPETITION

**D**evelopers, engineers, entrepreneurs, and innovators throughout the world are encouraged to form or join an AI XPRIZE team. To read more or submit an entry, visit [ai.xprize.org](http://ai.xprize.org).

experience, reason toward specific goals, and interact naturally with human collaborators. Thus, AI applications benefit greatly from expertise in data science and machine intelligence as well as more humanities-focused disciplines.

To foster increasing democratization and availability in AI technologies and allow individuals from many disciplines to participate in technological and social AI innovation, the competition is structured to attract large, collaborative teams. In fact, institutions and organizations with computing resources, datasets, tools, and domain expertise are encouraged to join the competition as supporting partners to help teams with the development and execution of their plans. Similarly, individual engineers, researchers, and specific domain experts (healthcare, environment, education, art, policy, economics, and so on) can join teams as advisors or hands-on support—indeed, it's this cross-pollination of knowledge, techniques, social contexts, and methodologies that will contribute to accelerating education and skills development as it drives groups toward meeting their goals and winning the prize.

**T**he XPRIZE's public visibility will have a tremendous impact on increasing our everyday awareness of AI, particularly the ways in which it benefits society. Through this kind of highly publicized, global, interdisciplinary, public/private competition, the XPRIZE Foundation also opens up a crucial forum for a variety

of stakeholders to help define the guidelines and mechanisms by which AI could or should be used. The benefits of this kind of heterogeneous, team-based approach to address important humanitarian challenges is that many voices can be heard, allowing us to more knowledgeably shape our future through technology.

Because this competition is, at its core, a large-scale collaboration employing interdisciplinary approaches to problem solving, it's a tremendous learning opportunity—indeed, the first real benefit is learning. 

*This article originally appeared in Computer, vol. 49, no. 11, 2016.*

**AMIR BANIFATEMI** is the lead for the XPRIZE Foundation's IBM Watson AI XPRIZE. Contact him at [amir.banifatemi@xprize.org](mailto:amir.banifatemi@xprize.org).

**JEAN-LUC GAUDIOT** is a professor of electrical engineering and computer science at University of California, Irvine. He is IEEE Computer Society's president-elect. Contact him at [gaudiot@uci.edu](mailto:gaudiot@uci.edu).

**myCS** Read your subscriptions through the myCS publications portal at <http://mycs.computer.org>



## There Is No AI Without IA

Seth Earley, *Earley Information Science*

**A**rtificial intelligence (AI) is increasingly hyped by vendors of all shapes and sizes—from well-funded startups to the well-known software brands. Financial organizations are building AI-driven investment advisors.<sup>1</sup> Chat bots provide everything from customer service<sup>2</sup> to sales assistance.<sup>3</sup> Although AI is receiving a lot of visibility, the fact that these technologies all require some element of knowledge engineering, information architecture, and high-quality data sources is not well known. Many vendors sidestep this question or claim that their algorithms operate on unstructured information sources, “understand” those sources, interpret the user query, and present the result without predefined architectures or human intervention. That very well might be true in certain circumstances, but most applications require a significant amount of hard work on the part of humans before neural nets, machine learning, and natural language processors can work their magic.

DigitalGenius, a product that received press coverage at a confer-

ence in 2015,<sup>4</sup> uses deep learning and neural nets. DigitalGenius first classifies the incoming questions into one of several categories for further processing: product information, account information, action request, comparison question, recommendation question, and so on. These classifications qualify as a foundational element of information architecture. The starting point is contextualizing the query, which is then passed to other modules, including product information systems and other databases and APIs. Each of these systems and sources in turn needs to be well architected to return the correct information. If the information is not structured and curated in some way, there is nothing for the system to return. The ability of DigitalGenius to use AI to engage customers is predicated on having high-quality, structured data.

### Digital Engagement: The Right Information in Context

Organizations are in a never-ending cycle of improving their digital means of engaging with

customers. These initiatives include improving personalization of the user experience by presenting more relevant content; tuning search results to return exactly what the user is interested in; and improving the effectiveness of offers and promotions. Organizations might also strive to increase response rates to email communications; provide better customer self-service; increase participation in user communities and other social media venues; and generally enhance the product experience through various other online mechanisms. In each of these cases, the means of engagement is providing a relevant piece of data or content (in the form of promotions, offers, sales, next-best action, products for cross sell and upsell, answers to questions, and so on) at exactly the right time and in the context that is most meaningful and valuable to the user.

This is done by interpreting the various signals that users provide through their current and past interactions with the organization. These signals include prior purchases, real-time click-stream

data, support center interactions, consumed content, preferences, buying characteristics, demographics, firmographics, social media information, and any other “electronic body language” that is captured by marketing automation and integration technologies. A search query could return a different result for a technical user than a nontechnical user, for example. At its core, search is a recommendation engine. The signal is the search phrase, and the recommendation is the result set. The more that is known about the user, the more the recommendation can be tailored. If the recommendation is related to a product, then clean, well-structured product data is a precondition.

### Personalization, User Signals, and Recommendations

Getting these recommendations right and truly personalizing the user experience requires that product data be correctly structured and organized, that content processes be integrated into product onboarding, and that associations can be made among products, content, and user intent signals. The relationship of products to content is based on knowledge of the user’s task and what content would help him or her accomplish it. The task might be a review, a how-to, product specifications, reference materials, instructions, diagrams and images, or other content that helps the user decide on the purchase.

AI encompasses a class of application that allows for easier interaction with computers and also allows computers to take on more of the types of problems that were typically in the realm of human cognition. Every AI program interfaces with information, and the better that information is structured, the more effective the

program is. A *corpus* of information contains the answers that the program is attempting to process and interpret. Structuring that information for retrieval is referred to as *knowledge engineering* and the structures are called *knowledge representation*.

### Ontologies as Knowledge Representations

Knowledge representation consists of taxonomies, controlled vocabularies, thesaurus structures, and all of the relationships between terms and concepts. These elements collectively make up the *ontology*. An ontology represents a domain of knowledge and the information architecture structures and mechanisms for accessing and retrieving answers in specific contexts. Ontologies

ent the right content for the user’s context, and requires not only that our customer data is clean, properly structured, and integrated across multiple systems and processes but also that the system understand the relationship between the user, his or her specific task, the product, and the content needed—all assembled dynamically in real time. Building these structures and relationships and harmonizing the architecture across the various back-end platforms and front-end systems results in an enterprise ontology that enables a personalized, omnichannel experience. Some might call this an enterprise information architecture; however, there is more to it than the data structures. Recall that the definition of an ontology includes real-world

**Every AI program interfaces with information, and the better that information is structured, the more effective the program is.**

can also capture “common sense” knowledge of objects, processes, materials, actions, events, and myriad more classes of real-world logical relationships. In this way, an ontology forms a foundation for computer reasoning even if the answer to a question is not explicitly contained in the corpus. The answer can be inferred from the facts, terms, and relationships in the ontology. In a practical sense, this can make the system more user friendly and forgiving when the user makes requests using phrase variations, and more capable when it encounters use cases that were not completely defined when it was developed. In effect, the system can “reason” and make logical deductions.

Correctly interpreting user signals enables the system to pres-

logic and relationships. The ontology can contain knowledge about processes, customer needs, and content relationships.<sup>5</sup>

### Mining Content for Product Relationships

Consumer and industrial products need to be associated with content and user context, but content can also be mined to suggest products for a user context. For an industrial application, a user might need parts and tools to complete maintenance on a hydraulic system. Using adaptive pattern-recognition software to mine reference manuals about hydraulic systems and repair, the system can extract a list of needed tools and related parts. A search on hydraulic repair will present a dynamically generated product

page based on the product relationships and correlated with the company's offerings. To some information professionals, this might sound complex and cumbersome to implement; however, there are emerging approaches that bring these aspirations closer to reality.

## If It Works, It's Not AI

The concept of what constitutes AI has evolved as technology has evolved. A colleague of mine has said, "It's artificial intelligence until you know how it works." An interesting perspective indeed. I found support for this in material from an MIT AI course (see <http://bit.ly/1WTCUXK>):

*There's another part of AI [...] that's fundamentally about applications. Some of these applications you might not want to call "intelligent" [...] For instance, compilers used to be considered AI, because [...] statements [were in a] high-level language; and how could a computer possibly understand? [The] work to make a computer understand [...] was taken to be AI. Now [...] we understand compilers, and there's a theory of how to build compilers [...] well, it's not AI anymore. [...] When they finally get something working, it gets co-opted by some other part of the field. So, by definition, no AI ever works; if it works, it's not AI.*

Seemingly intractable problems have been solved by advances in processing power and capabilities. Not long ago, autonomous vehicles were considered technologically infeasible due to the volume of data that needed to be processed in real time. Speech recognition was unreliable and required extensive speaker-dependent training sessions. Mobile phones were once "auto-mobile" phones, requiring a car trunk full of equipment (my

first car phone in the '80s, in addition to costing several thousand dollars, left little room for anything else in the spacious trunk). Most AI is quietly taken for granted today. The word processor I am using was once considered an advanced AI application!

## Simplicity Is Hidden Complexity

Under the covers, AI is complex; however, this complexity is hidden from the user and is, in fact, the enabler of an easy, intuitive experience. It is not magic, and it requires foundational structures that can be reused across many different processes, departments, and applications. These structures are generally first developed in silos and standalone tools; however, the true power will be realized when considered in a holistic framework of machine-intelligence-enabled infrastructure. AI will change the business landscape but will require investments in product and content architecture, customer data, and analytics, and the harmonization of tools in the customer engagement ecosystem. The organizations that adopt these approaches will gain significant advantages over the competition.

## Clean Data Is the Price of Admission

AI approaches are proposed as the answer to enterprise challenges around improving customer engagement by dealing with information overload. Before those approaches can be leveraged, however, organizations need to possess the data needed as inputs to machine learning algorithms that can in turn process these diverse signals from unstructured and structured sources. Clean, well-structured, managed data is assumed.

In many cases, the data being processed or corpus being analyzed by AI systems is typically less structured than better-organized sources such as financial and transactional data. Learning algorithms can be used to both extract meaning from ambiguous queries and attempt to make sense of unstructured data inputs. Humans might phrase questions using different terminology, or they might ask overly broad questions. They are not always clear about their objectives—they don't necessarily know what they are looking for. This is why human sales people typically engage prospects in conversations about their overall needs, rather than asking them what they are looking for. (At least, the good ones do.)

Inserting AI into the process is more effective when users know what they want and can clearly articulate it, and where there is a relatively straightforward answer. The algorithm does what it does best in terms of handling variations in how those questions are asked, interpreting the meaning of the question, and processing other unstructured signals that help further contextualize the user's intent. There are many flavors of AI and many classes of algorithm that comprise AI systems. However, even when AI systems are used to find structure from completely unstructured information, they still require structure at the data layer.

Given that the data being searched by the AI system is unstructured, why do we need information architecture? Unstructured information is typically in the form of text from pages, documents, comments, surveys, social media, or some other source. Though it is unstructured, there are still parameters associated with

**Table 1. Example applications for AI technology.**

<b>Application/problem</b>	<b>Typical approach</b>	<b>Why it will not work</b>	<b>How AI helps</b>
Omnichannel selling	Integrations across systems that are inflexible	Content is typically out of synch, and systems have brittle integrations.	Harvesting signals of user intent across channels and devices; intelligent integration across diverse classes of customer-engagement technologies.
Personalization and recommendation engines	Classification of users according to personas and presentation of content based on use cases	Assumptions are rudimentary and simplistic, limit user choice, and rarely anticipate needs.	Incorporating multiple data sources, ontologies, and reasoning algorithms.
Dynamic content assembly	Limited assembly based on curated content	The number of permutations cannot be tracked, managed, or understood.	Assembling content that meets a user's preferences while incorporating real-time data, such as news and market performance supported by knowledge embedded in ontologies.
Field service equipment maintenance	View of maintenance schedules and analysis of field reports.	The volume of information from sensors and hidden signals from degrading performance is not perceivable by humans.	Processing performance data to predict or prevent maintenance events.
Customer self-service	Knowledge bases and FAQs	There are insufficient considerations of user context, language variations, fast-changing content, and a high volume of possible use cases.	Enabling intelligent agent knowledge base access using multiple variables to provide specific answers with metrics-driven feedback loops.
Call center support	Reps trained through on-the-job experience, knowledge bases	For technical and specialized knowledge, the level of experience needed is too costly.	Enabling intelligent agent access integrating answers and content from multiple knowledge and data sources using natural language interfaces.
Product and tool presentation	Hand-built landing pages	Maintenance of hand-curated and assembled landing pages becomes costly and unmanageable.	Harvesting product relationships through knowledge acquisition and mining of reference materials based on maintenance tasks and repair use cases.

the source and context. Social media information requires various parameters to describe users, their posts, relationships, time and location of posts, links, hashtags, and so on. The information architecture question in this case characterizes the structures of the input data, so that the system can be programmed to find patterns of interest. Even in the case of unsupervised machine learning (a class of application that derives signals from data that are not predefined by a human), the programmer still needs to describe the data in the first place with attributes and values. There might not be predefined categories for outliers and patterns that are

identified, but the input needs to have structure.

A common fallacy in considering big data sources that form inputs for machine learning is that because the data is “schema-less” (does not have a predefined structure), no structure is required. Data still requires attribute definitions, normalization, and cleansing to apply machine learning and pattern-identification algorithms.<sup>6</sup> As enterprises embark on the path to machine learning and AI, they should, first and foremost, be developing an enterprise ontology that represents all of the knowledge that any AI system they deploy would process, analyze, leverage, or require.

Some vendors might debate the value of this approach, insisting that their algorithms can handle whatever you throw at them—however, I would argue that this is only the case when the ontologies are self-contained in the tool. Even so, there will always be gaps between what a tool developed for broad adoption can contain and the specialized needs of an enterprise. Even if the tool is developed for a specific industry, the differences in processes found from organization to organization require specialized vocabularies and contextual knowledge relationships. This is a significant undertaking; however, not doing so misses an essential step in the process.

**Table 2. Examples of AI tools with representative applications, limitations, considerations, and data sources.**

<b>Class of tool</b>	<b>Best suited for</b>	<b>Limitations</b>
Inference engine	Deriving product and data relationships from unstructured content	Significant effort required in certain cases for development of adaptive pattern-recognition algorithms
Intelligent agent	Highly selective search and information retrieval for definable processes	Breadth of possible use cases, variability of queries and questions, degree of manual term-relationship mapping
Auto-classifier	Large volumes of higher-quality content where training sets or clear rules can be applied	Rules bases can become very complex, high volumes of training content required, variability of content
Entity extractors	Predictable formats for data (Social Security numbers, addresses, names, phone numbers, account numbers)	Ambiguity of entity values (for “Washington,” can be Washington State, George Washington, Washington, DC, and so on), variations in format, quality of content
Unsupervised machine learning	Pattern detection, recognition, and prediction; anomaly detection, outliers, hidden attributes, and relationships; discovering new patterns or segmenting audiences, content, or data into clusters and groupings	Depends on creation of foundational hypothesis to define type of outlier or pattern (though not necessarily details of pattern); selection of class of algorithm requires technical sophistication and iterative testing across many classes of unsupervised learning from data signals to content, images, or transactions
Supervised machine learning	Pattern detection based on training data and used for well-understood patterns; discovering data, content, and relationships based on examples; finding similar documents, purchase patterns, audiences with like characteristics; predicting outcomes	Requires training sets and example data that teach the algorithm what type of information is being sought; needs large enough samples to test various hypotheses and an understanding of the specific outcome sought; risk of finding patterns where none exist (over fitting)
Hybrid unsupervised and supervised learning	Finding hidden patterns (unsupervised); using those patterns to train an algorithm to locate more instances of the data or content	Similar in concept to creating a survey with open-ended questions, then taking the answers to the survey and using those results to form closed-ended surveys

Much of what is described as AI is an extension of well-known approaches to addressing information management problems, all of which require clean foundational data and information structures as a starting point. The difference between standard information management and practical AI lies in understanding the limits of these technologies and where they can best be applied to address enterprise challenges.

The remainder of this article describes how your organization can identify use cases that will benefit from AI, identifies data sources that offer reliable and meaningful insights to train and guide AI, and defines governance, curation, and scalable processes

that will allow for continuous improvement of AI and cognitive computing systems.

### Identifying Use Cases

Differentiating AI use cases from standard information management use cases requires considering the sources of data that comprise the “signals” being processed, the type of task the user faces, and the systems that will be part of the solution. The difference in approaches to these problems lies in how the sources of data will be curated and ingested, how organizing principles need to be derived and applied, the sophistication of the functionality desired, and the limitations of the current solution in place. An

AI approach will require a greater level of investment, sponsorship at the executive level, program-level governance, and an enterprise span of influence. It will also require a longer-term commitment than a typical information management project. Although there are opportunities to deploy limited scope AI, fully leveraging it as a transformative class of technology should be part of an overall digital transformation strategy in some cases on the scale of enterprise resource planning (ERP) programs, with commensurate support, funding, and commitment. (Some ERP programs can cost US\$50 to \$100 million or more). Although no organization will approach an unproven set of technologies with

Decision factors	Example data
Volume of content, quality of corpus, structure of data within content (use of tables, headers, other identifiers, and delimiters)	Structured documents such as catalogs, engineering specifications, knowledge bases, ERP data
Narrow scenarios, defined repeatable tasks, sophisticated audience, standard terminology for domain	Policy and procedure manuals, chat logs, customer service databases, support tickets, knowledge harvested from human subject matter experts
High-value content that has a predictable structure, curated with editorial standards	Semi-structured content such as regulations, journal articles, news articles, policies and procedures, support documentation, engineering documents
Security, compliance, and reference-based applications (looking for personally identifiable information, sensitive account information, intellectual property protection, fact identification)	Same semi-structured content sources mentioned above, along with controlled vocabularies and thesaurus structures to identify entity values
Determination of type of patterns that are being detected; not typically selected in isolation because learning algorithms are embedded in many aspects of AI application; requires guidance of data science specialist knowledgeable in specific type of algorithm	Unsupervised learning can be applied to virtually any data source as the tools look for patterns; can be applied to unstructured and semi-structured content as well as transactional data, images, audio and video, scientific datasets, sensor data, social media sources, and so on
Similar challenges to unsupervised learning: typically embedded in AI applications rather than being applied as standalone tool and requires specialized knowledge; can be used to find types of images, content or data, purchase patterns, audiences, users, combinations of products, or other patterns depending on data sources	Same as unsupervised; the difference is having training sets that contain patterns sought
Most AI programs contain hybrid implementations of supervised and unsupervised and therefore the complexity of each approach can be compounded; requires research and experimentation to find correct combinations and fine-tune algorithms	Same data used for unsupervised learning can be used as input for supervised learning after representative patterns have been identified

that level of commitment, funding needs to be allocated to extend proven approaches with emerging AI technologies.

The roadmap for AI transformation includes continuous evaluation of payback and ROI and focuses on short-term wins while pursuing longer-term goals. Most organizations are attempting to solve the problems described in Table 1 with limited approaches, departmental-level solutions, standalone tools, and insufficient funding. These problem classes are facing most enterprises, and though progress can be made with limited resources and siloed approaches, this would be an extension of business as usual. Truly transformative applications will require an enterprise view of the

organization’s knowledge landscape and implementation of new governance, metrics, and data quality programs—governance to make decisions, metrics to monitor the effectiveness of those decisions, and data quality to fuel the AI engine.

Table 1 presents example applications for AI technology.

### Identifying Data Sources

Training data can come from typical knowledge bases, the more highly curated the better. Call center recordings and chat logs can be mined for content and data relationships as well as answers to questions. Streaming sensor data can be correlated with historical maintenance records, and search logs can be mined for use cases and

user problems. Customer account data and purchase history can be processed to look for similarities in buyers and predict responses to offers; email response metrics can be processed with text content of offers to surface buyer segments. Product catalogs and data sheets are sources of attributes and attribute values. Public references can be used for procedures, tool lists, and product associations. YouTube video content audio tracks can be converted to text and mined for product associations. User website behaviors can be correlated with offers and dynamic content. Sentiment analysis, user-generated content, social graph data, and other external data sources can all be mined and recombined to yield knowledge and

user-intent signals. The correct data sources depend on the application, use cases, and objectives.

Table 2 describes examples of AI tools with representative applications, limitations, considerations, and data sources. Though not meant to be an exhaustive list, and recognizing that one class of tool is frequently leveraged in other tools and applications (an intelligent agent can use inference engines, which in turn can leverage learning algorithms, for example), the table articulates considerations for exploring one approach versus another.

## Defining Governance, Curation, and Scalable Processes

AI and cognitive computing are managed in the same way as many other information and technology governance programs. They require executive sponsorship, charters, roles and responsibilities, decision-making protocols, escalation processes, defined agendas, and linkage to specific business objectives and processes. These initiatives are a subset of digital transformation and are linked to customer life cycles and internal value chains. Because the objective is always to affect a process outcome, all AI and cognitive computing programs are closely aligned with ongoing metrics at multiple levels of detail—from content and data quality to process effectiveness and satisfaction of business imperatives—and ultimately are linked to the organizational competitive and market strategy. Milestones and stages are defined to release funding for program phases, each with defined success criteria and measurable outcomes.

**A**I will no doubt continue to impact every aspect of our personal and professional lives. Much of this impact

will occur in subtle ways—such as improved usability of applications and findability of information. These will not necessarily appear on the surface to be AI. Over time, AI-driven intelligent virtual assistants will become more fluent and capable, and will become the preferred mechanism for interacting with technology. Humans create knowledge, while machines process, store, and act on that knowledge. AI is applied human knowledge. Organizations need to develop the foundations for advancing AI by capturing and curating that knowledge and by building the foundational data structures that form the scaffolding for that knowledge. Without those components, the algorithms have nothing to run on. 

## References

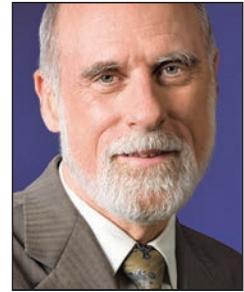
1. J. Vögeli, “UBS Turns to Artificial Intelligence to Advise Clients,” Bloomberg, 7 Dec. 2014; [www.bloomberg.com/news/articles/2014-12-07/ubs-turns-to-artificial-intelligence-to-advise-wealthy-clients](http://www.bloomberg.com/news/articles/2014-12-07/ubs-turns-to-artificial-intelligence-to-advise-wealthy-clients).
2. C. Green, “Is Artificial Intelligence the Future of Customer Service?” MyCustomer, 3 Dec. 2015; [www.mycustomer.com/service/channels/is-artificial-intelligence-the-future-of-customer-service](http://www.mycustomer.com/service/channels/is-artificial-intelligence-the-future-of-customer-service).
3. E. Dwozkin, “Can Artificial Intelligence Sell Shoes?” blog, *WallStreetJ*, 17 Nov. 2015; <http://blogs.wsj.com/digits/2015/11/17/can-artificial-intelligence-sell-shoes/>.
4. R. Miller, “DigitalGenius Brings Artificial Intelligence to Customer Service via SMS,” Tech Crunch, 5 May 2015; <http://techcrunch.com/2015/05/05/digitalgenius-brings-artificial-intelligence-to-customer-service-via-sms/>.
5. S. Earley, “Lessons from Alexa: Artificial Intelligence and Machine Learning Use Cases,” blog, Earley Information Science, 24 Mar. 2016; [www.earley.com/blog/lessons-alex-artificial-intelligence-and-machine-learning-use-cases](http://www.earley.com/blog/lessons-alex-artificial-intelligence-and-machine-learning-use-cases).
6. J. Brownlee, “How to Prepare Data for Machine Learning,” Machine Learning Mastery, 25 Dec. 2013; <http://machinelearningmastery.com/how-to-prepare-data-for-machine-learning/>.

*Seth Earley is CEO of Earley Information Science ([www.earley.com](http://www.earley.com)). He’s an expert in knowledge processes, enterprise data architecture, and customer experience management strategies. His interests include customer experience analytics, knowledge management, structured and unstructured data systems and strategy, and machine learning. Contact him at [seth@earley.com](mailto:seth@earley.com).*

*This article originally appeared in IT Professional, vol. 18, no. 3, 2016.*



# Unreliable Machine Learning



Vinton G. Cerf • Google

I'm not an expert in machine learning, so this essay is likely to expose this fact. Machine learning has been in the news increasingly in the past few years with the occasional spectacular headline, such as "AlphaGo Beats World Champion Go Player!" or "Self-Driving Cars Pass Two Million Miles on the Road!" We can't help being impressed by the successes of neural networks and their uncanny ability to perform tasks that in some cases surpass human capacity. So how does this work and why am I worried about it?

John Giannandrea, Google's Senior Vice President for Search, Research, and Machine Intelligence, likens neural networks to a box with about a million dials on it. As the box is exposed to inputs, it makes choices or decisions, and is given feedback as to the correctness or desirability of its output. As these lessons proceed, the box adjusts the dials until it performs to the trainer's satisfaction. On the other hand, if someone said, "What happens if you adjust THIS dial by 2 percent?" it's not clear that anyone (including the trainer) would have a credible answer. This is, of course, where my general lack of knowledge about these mechanisms presumably shows. My impression is that if such devices are put into operation and make a poor decision under some set of conditions, the only option is to analyze the conditions leading to the bad choice and to invent new training sequences to get the system to magically adjust its dials for that case. After which, someone then needs to test to see that other inputs still produce satisfactory results.

There's something mildly unsettling about this – it seems to suggest that we don't have a theory of operation that allows us to make direct adjustments in a confident way. When in doubt, Google! So I did, and discovered a very cogent paper from nearly 20 years ago on the subject of neural networks and Bayesian methods. My interpretation of this paper, as applied to my worry about theory of operation, is that the Bayesian methods can be applied to the operation of the neural network to corral some of the side effects of overfitting in consequence of limited input data. While this doesn't offer a mechanism for direct adjustment of the neural model parameters, short of training, it does seem to improve the neural network's performance. Presumably, some progress has been made in the past 20 years so that Bishop's paper doesn't represent the state of the art, although I found it oddly comforting.

If we presume that neural networks will find their way into an increasing variety of processes, we might wonder how these systems will behave when they interact with each other in the real world. Reading the headlines again, my impression is that there are a number of companies out to invent self-driving cars, Google included. It would be wrong to assume that these cars rely solely on neural methods for their operation. A great deal of hard information is available for navigation, state information (traffic lights, for example), detection of moving and fixed objects, collision avoidance, and so on. If there's a significant neural component

to the operation of these cars, we might wonder how they will interact with each other.

In the stock market, programmed trading is now a major component of all the trades done and multiple parties run these systems. I've been told that they sometimes behave like young kids playing soccer: they all run after the ball, leading to various kinds of instability. I wonder about streets filled with self-driving cars, each running its own brand of neural network. Each will have been trained with different scenarios and each might have its own "blind

spots" with regard to particular situations. Will the training feedback be based on accidents? "Back to driver training for you, Robot 45792B!"

**A**s with all my essays, I hope that I provoke some discussion so that I can learn, as well. Of course, I might provoke outrage among those who know a lot more than I do on this topic, but that's the risk you take in writing for public consumption. ☐

**Vinton G. Cerf** is vice president and chief Internet evangelist at Google, and past

president of ACM. He's widely known as one of the "fathers of the Internet." He's a fellow of IEEE and ACM. Contact him at [vgcerf@gmail.com](mailto:vgcerf@gmail.com).

*This article originally appeared in IEEE Internet Computing, vol. 21, no. 1, 2017.*

myCS

Read your subscriptions through the myCS publications portal at <http://mycs.computer.org>.

IEEE  computer society

Read all your IEEE magazines and journals yourWAY on

myCS

Introducing **myCS**, the digital magazine portal from IEEE Computer Society. Go beyond static, hard-to-read PDFs with an easily accessible, customizable, and adaptive experience.

**There's No Additional Cost!**

Now there's even more to love about your membership...

**LEARN MORE AT: [mycs.computer.org](http://mycs.computer.org)**

IEEE  computer society

# ROCK STARS ADAPTIVE CYBERSECURITY

June 6, 2017 | Waltham, MA

## CYBERSECURITY ATTACKS - HOW WILL YOU RESPOND?

The landscape is ever-changing.  
Find out how to win the cyber  
war at Rock Stars of Adaptive  
Cybersecurity.

[www.computer.org/AdaptiveCyber](http://www.computer.org/AdaptiveCyber)



**Caleb Barlow**  
V.P. Threat  
Intelligence,  
IBM Security



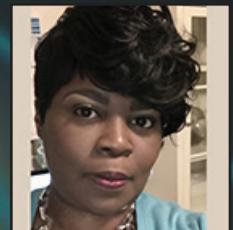
**Edna Conway**  
CSO, Global  
Value Chain,  
Cisco



**Jim Routh**  
CSO, Aetna



**Gary Hayslip**  
CISO, Deputy  
Director,  
City of San Diego



**Shon Lyublanovits**  
IT Security Subcategory  
Manager, U.S. General  
Services Administration



# Masterminds of Artificial Intelligence

## Marvin Minsky and Seymour Papert

**George Strawn**, *US National Academies of Sciences, Engineering, and Medicine*  
**Candace Strawn**

**T**he world lost two great scientists this year. Marvin Minsky passed away on 24 January 2016, and Seymour Papert passed away on 31 July. They were both 88 years old. Their careers at MIT were closely joined, both in artificial intelligence (AI) and in other fields. This article will highlight some of their contributions and one “controversy” that they jointly sparked. There are a number of YouTube videos of talks and lectures given by both of these masterminds that are well worth watching (see, for example, <https://youtu.be/-pb3z2w9gDg>). Moreover, a friend’s chapter-length biography of Minsky recently appeared in Stephen Wolfram’s *Idea Makers*.<sup>1</sup>

### Early Life and Education

Minsky was born in 1927 in New York City to prosperous parents. He received a good education from the beginning, including at the

Bronx High School of Science and the Phillips Academy. By the last years of World War II, he was old enough to be drafted and spent 1944–1945 in the US Navy. At the end of hostilities, he returned to his schooling, obtaining a bachelor’s degree from Harvard in 1950 and a PhD from Princeton in 1954, both in mathematics. His thesis was titled “Theory of Neural-Analog Reinforcement Systems and Its Application to the Brain Model Problem.” This was an early contribution to the field of *artificial neural networks* (ANNs) which is the same field in which he and Papert created the aforementioned controversy, which we discuss later. He joined the MIT faculty in 1958 (see <http://web.media.mit.edu/~minsky/minskybiog.html>).

Papert was born in 1927 in South Africa and was educated there at the University of the Witwatersrand, receiving a BA in philosophy in 1949 and a PhD in mathematics in

1952 (see [https://en.m.wikipedia.org/wiki/Seymour\\_Papert](https://en.m.wikipedia.org/wiki/Seymour_Papert)). He then moved to England, where he worked at St. John’s College, Cambridge, and earned a second PhD in mathematics in 1959. Papert was a prominent activist against South African apartheid policies during his university education. He also worked at the Henri Poincaré Institute at the University of Paris, the University of Geneva (where he worked with Jean Piaget; see [https://en.m.wikipedia.org/wiki/Jean\\_Piaget](https://en.m.wikipedia.org/wiki/Jean_Piaget)), and the National Physical Laboratory in London before becoming a research associate at MIT in 1963.

### Careers at MIT

Minsky is best known as a father of AI, having cofounded the MIT AI laboratory in 1959 (with John McCarthy, who subsequently moved to Stanford). However, he made contributions in many domains, including graphics,

symbolic mathematical computation, knowledge representation, commonsensical semantics, machine perception, and both symbolic and connectionist learning. He was also a pioneer of robotics and telepresence. He designed and built some of the first visual scanners—mechanical hands with tactile sensors. He also built one of the first “turtles” for Papert’s LOGO programming language, including the software and hardware interfaces. He published a major work summarizing many years of thought in 1985, titled *The Society of Mind* ([https://en.m.wikipedia.org/wiki/Society\\_of\\_Mind](https://en.m.wikipedia.org/wiki/Society_of_Mind)).<sup>2</sup>

Papert is perhaps best known for his contributions to education, and he also made contributions in other fields. Minsky appointed him codirector of the MIT AI laboratory in 1967. Based on Piaget’s constructivist learning theory, for which he was a principle advocate, Papert codesigned the LOGO programming language, which influenced later languages such as Smalltalk and Scratch (see [https://en.m.wikipedia.org/wiki/Logo\\_\(programming\\_language\)](https://en.m.wikipedia.org/wiki/Logo_(programming_language))). In 1980, he published *Mindstorms: Children, Computers, and Powerful Ideas* ([https://en.m.wikipedia.org/wiki/Mindstorms\\_\(book\)](https://en.m.wikipedia.org/wiki/Mindstorms_(book))).<sup>3</sup>

### Minsky’s Society of Mind

*The Society of Mind* is more speculative cognitive science and philosophy than AI. That is, it proposes that the human mind is a “society” of independent agents that evolution produced over time to deal with various problems that humans needed to solve. The book consists of 270 one-page chapters, each dealing with a brain agent, process, or interconnection issue.

In 2006, Minsky published a sequel, *The Emotion Machine*,<sup>4</sup> which proposes theories that could account for human higher-level

feelings, goals, emotions, and conscious thoughts in terms of multiple levels of processes, some of which can reflect on the others. This view of consciousness is process-oriented, in that consciousness is seen to be the result of more than 20 processes going on in the human brain.

If *The Society of Mind* and *The Emotion Machine* contain ideas that can be implemented in computers, then Minsky has indeed provided a path for AI researchers to follow. If cognitive neuroscience discovers that the agents and processes he identifies are indeed present in the structures of the brain, then he has identified major components of “natural intelligence.” Perhaps the time is not too far off when these conjectures will be supported or refuted.

### Papert’s LOGO

As mentioned, Papert had worked with Piaget in Geneva and was an advocate for his constructivist learning theory ([https://en.m.wikipedia.org/wiki/Constructivism\\_\(philosophy\\_of\\_education\)](https://en.m.wikipedia.org/wiki/Constructivism_(philosophy_of_education))). Papert’s ideas were given the confusingly similar name of constructionist learning theory. In a US National Science Foundation proposal, he defined this theory as follows (see [https://en.m.wikipedia.org/wiki/Constructionism\\_\(learning\\_theory\)](https://en.m.wikipedia.org/wiki/Constructionism_(learning_theory))):

*Constructionism is a mnemonic for two aspects of the theory of science education underlying this project. From constructivist theories of psychology, we take a view of learning as a reconstruction rather than as a transmission of knowledge. Then we extend the idea of manipulative materials to the idea that learning is most effective when part of an activity the learner experiences as constructing a meaningful product.*

The idea of the (child) learner constructing a meaningful prod-

uct was realized by constructing LOGO programs and robotic devices that could be controlled by LOGO programs. Papert’s coinventor of LOGO, Wally Feurzeig ([https://en.m.wikipedia.org/wiki/Wally\\_Feurzeig](https://en.m.wikipedia.org/wiki/Wally_Feurzeig)), had been developing conversational programming languages at BBN based on Fortran, which was designed for mathematical problems. By focusing on nonmathematical problems, they defined a language simple enough for children to learn. And by connecting the language to the control of “turtles” and other robotic devices, they provided additional motivation (fun) for young learners.

Papert expressed these ideas in *Mindstorms: Children, Computers, and Powerful Ideas*.<sup>3</sup> He also collaborated with the construction toy manufacturer Lego on their LOGO-programmable Lego Mindstorms robotics kits, which were named after his book. The third generation of Lego Mindstorms was released in 2013 (see [https://en.m.wikipedia.org/wiki/Lego\\_Mindstorms\\_EV3](https://en.m.wikipedia.org/wiki/Lego_Mindstorms_EV3)). A Lego Mindstorms kit contains software and hardware to create customizable, programmable robots. It includes an intelligent brick computer that controls the system, a set of modular sensors and motors, and Lego parts from the Technic line to create the mechanical systems. (Several different programming languages can now be used in Lego Mindstorms.)

### Artificial Neural Net Controversy and Deep Learning

The subject of ANNs goes back to 1943, when Warren McCulloch and Walter Pitts created a computational model for neural networks ([https://en.m.wikipedia.org/wiki/Artificial\\_neural\\_network](https://en.m.wikipedia.org/wiki/Artificial_neural_network)). As mentioned, Minsky’s PhD dissertation had contributed to the early study of ANNs. But in

1969, Minsky and Papert coauthored a book titled *Perceptrons*,<sup>5</sup> which gave rise to a controversy (over what they really said) that is still being argued today. In their book, they showed that “simple ANNs” were quite limited in computational power. The assumption that “complex ANNs” (technically called hidden-layer ANNs) were also of limited power might have been made by Minsky and Papert, but not in this book. One commentator stated that “the often-miscited Minsky/Papert text caused a significant decline in interest and funding of neural network research,” and that it would take 10 years for neural network research to experience a resurgence in the 1980s. (An expanded edition of *Perceptrons*<sup>6</sup> was reprinted in 1987 with some errors from the original text shown and corrected; see <https://en.m.wikipedia.org/wiki/Perceptron>). Another commentator, a researcher at the MIT AI lab, quoted Minsky and Papert in the 1971 “Report of Project MAC,”—directed at funding agencies—on “gamba networks” (see [https://en.m.wikipedia.org/wiki/Perceptrons\\_\(book\)](https://en.m.wikipedia.org/wiki/Perceptrons_(book))):

*Virtually nothing is known about the computational capabilities of [these machines]. We believe that [they] can do little more than can a low order perceptron.*

(See [www.ucs.louisiana.edu/~isb9112/dept/phil341/histconn.html](http://www.ucs.louisiana.edu/~isb9112/dept/phil341/histconn.html) for a second view of this history.)

By the 2000s, hidden-layer ANNs had proven to be very powerful and in fact were powering many real AI applications. Hidden-layer ANNs are now part of the family of deep learning algorithms ([https://en.m.wikipedia.org/wiki/Deep\\_learning](https://en.m.wikipedia.org/wiki/Deep_learning)), which has propelled us from an “AI winter” into an “AI summer.”

**T**he authority of Minsky and Papert was evident in that their negative view of ANNs was probably a factor in discouraging researchers and funders of research from further study of ANNs at that time. Their authority arose, of course, from their many important contributions to AI and allied fields. They will be remembered more for those contributions than for the controversy. **IT**

## References

1. S. Wolfram, *Idea Makers: Personal Perspectives on the Lives & Ideas of Some Notable People*, Wolfram Media, 2016.
2. M. Minsky, *The Society of Mind*, Simon & Schuster, 1988.

3. S.A. Papert, *Mindstorms: Children, Computers, and Powerful Ideas*, Basic Books, 1980.
4. M. Minsky, *The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind*, Simon & Schuster, 2006.
5. M. Minsky and S.A. Papert, *Perceptrons*, MIT Press, 1969.
6. M. Minsky and S.A. Papert, *Perceptrons: An Introduction to Computational Geometry*, expanded ed., MIT Press, 1987.

**George Strawn** is the director of the Board for Research Data and Information at the US National Academies of Sciences, Engineering, and Medicine. He is the former director of the National Coordination Office for the Networking and Information Technology Research and Development Program (NITRD). Contact him at [gostrawn@gmail.com](mailto:gostrawn@gmail.com).

**Candace Strawn** is a retired high school, community college, and university teacher. Contact her at [castrawn@gmail.com](mailto:castrawn@gmail.com).

This article originally appeared in *IT Professional*, vol. 18, no. 6, 2016.







This series of in-depth interviews with prominent security experts features Gary McGraw as anchor. *IEEE Security & Privacy* magazine publishes excerpts of the 20-minute conversations in article format each issue.

[www.computer.org/silverbullet](http://www.computer.org/silverbullet)  
 \*Also available at iTunes



## Risks

Matthew Turk | National Center for Supercomputing Applications

**M**y thoughts about the topic of this department are usually in a project context—for me, “scientific programming” typically takes the form of something related to computational physics, or data analysis and visualization. I proceed through a few stages when evaluating a new or existing project: deciding what needs to happen, sketching out a rough plan, evaluating my options, and determining the manageable risks. Particularly when preparing open source scientific projects or those that consume large quantities of resources, risk evaluation is crucial, both to commit to a project and to harden it against possible weaknesses.

I’ve spent a fair amount of time evaluating how I think about risk and how it has influenced my projects and plans. As I’ve tried to be more self-aware about this process (especially considering that “being aware of risk” is often a nice way of saying “being a tad bit afraid of it”), I’ve also realized that implicit in the decision-making process is risk mitigation.

If you can’t mitigate risks out of existence, you still have to deal with—and attempt to counter—them in some fashion.

Three particular risk scenarios have been on my mind recently, and I have a few strategies for addressing them. These won’t work for everyone—they don’t even always work for me—but they could be of some use, depending on the situation.

### Risk of Making a Mistake

The risk that I worry about the most, the one that keeps me up at night, is the risk of making a mistake. The dangers are both overt and subtle. The vagaries of transcription, the translation of an algorithm into a programming language, for example, can result in changes in functionality—what if there’s a stray parenthesis, a missing decimal point, or something even more insidious, such as a line of code extending past the character limit in a Fortran routine? Fortunately, an algorithm isn’t the same as an implementation; for all our lofty goals of translating formulas into machine code, they still remain separate. But beyond minor

Perfection can't be reached; this must be taken as a given. Instead, we should focus on mistaken identification, mitigation, and reduction. How can we find mistakes in software? And, from the perspective of mitigating risks going forward, how can we guard against adding new mistakes?

mechanical errors, larger errors still lurk—perhaps a method doesn't converge with sufficiently low error, terms can't be split in some regimes, or a corner case has gone unnoticed. These mistakes, too, are risks. I've heard far too often from colleagues that they're apprehensive of releasing code for fear of exposing mistakes. The implicit irony is that they aren't apprehensive about utilizing their software for their own purposes but of exposure.

I'm reluctant to attribute this quotation, as I'm sure I'll misattribute it, but someone once said, "Are you afraid of releasing your software because it might have bugs in it? I'll save you some time. It does." This really stuck with me, because it says not only that this is to be expected and you aren't alone, but also that it's something you shouldn't be in denial about. The most self-aware expression of this that I've seen is the version numbering system used in TeX, which will converge to pi during the lifetime of its inventor, Donald Knuth.

The most obvious answer, of course, is to test things, which is relatively straightforward when the code being tested is modular, has few execution paths, and conforms to a result that can be identified a priori. But when an algorithm, module, or project doesn't possess these attributes, ever-increasing circles around them can still be identified and tested. You start by breaking the system down into its smallest parts, testing them in multiple combinations, and then testing those combinations. You need to search input space parameters and guard against unknown behavior—in contrast with a silent, implicit failure that can be misinterpreted as a result, an explicit failure when applied outside of known input space can save confusion and heartache. When explicit, analytic, or prior solutions can't be determined in advance, we often have to resort to testing against regressions through so-called "answer testing": if an answer is agreed on as being a gold standard, subsequent iterations of development shouldn't deviate from it. For more on this, see P.F. Dubois, "Testing Scientific Programs," *Computing in Science & Eng.*, vol. 14, no. 4, 2012, pp. 69–73, doi:10.1109/MCSE.2012.84; and K. Hinsén, "Technical Debt in Computational Science," *Computing in Science & Eng.*, vol. 17, no. 6, 2015, pp. 103–107, doi:10.1109/MCSE.2015.113.

The best way to mitigate—but not eliminate—risk is to openly share, discuss, and examine written code. Sharing code doesn't immediately and magically ensure that it's error free (it doesn't even guarantee that it will monotonically decrease in error count), but it does increase the opportunity for discovering mistakes.

### Risk of Depending on Something

The rate of change in the scientific software ecosystem can be very difficult to predict. Libraries interoperate through APIs that provide clearly defined relationships for what serves as input (and what in return serves as output from functions and routines). Although these APIs often remain stable, they're fundamentally outside the control of downstream end users—in this case, scientific programmers using those APIs. This type of dependency brings with it its own risks, with the software developer required to maintain and keep up with changes (or forego gaining access to improvements and bug fixes) as well as ensure that any behavioral changes are understood and accounted for.

Utilizing an external package or project requires giving up a degree of autonomy and self-direction. The converse, of course, is that simply implementing every item necessary requires an incredible amount of dedication and hard work that could be much more easily spent on improving the specific project being worked on.

Mitigating this risk is in many ways more difficult than mitigating the risk of errors, as it involves people rather than code. I've found that the most productive way of mitigating the risks in adding a dependency, building in an ecosystem, or interoperating with a tool is to understand the alignment of interests. Are the motivations of the external library or software developers aligned with my own? Or at the very least, are they not antialigned? What's on my road map, what's on theirs, and do they allow for long-term interoperability?

For many pieces of software, this is a simple evaluation: API contracts are stable unless specifically noted, and in some fields of scientific software, these notes come once a decade. But in other areas, particularly the parts of the ecosystem undergoing rapid change (or that are well-suited to rapid change, such as dynamic languages or anything related to Web services), these situations must be more carefully monitored. This can take the form of extensive integration tests between systems, "pinning" specific version numbers as dependencies, and building in contingency plans for if a piece of software shifts in focus or no longer provides the necessary functionality.

### Risk of Starting Something Big

Being wrong and building on an unreliable foundation are equally scary, but they both have mitigation paths that are

## These are the easy parts to implement; best practices of software development help guide these types of invasive projects and changes. But how do we convince ourselves that this is the right path?

easy to see and straightforward to implement. But there's another risk that lurks in the minds of scientific programmers, and that's the risk of starting something all-caps BIG. Major undertakings, refactorings, implementations—these are the investments of time, resources, and energy that might or might not flourish. If your fantastic new project idea results in enormous benefits, that's marvelous, but what if it results in the alienation of collaborators? If your project brings in new people, ideas, and results that are outstanding and exciting, it's a success, but how do you mitigate against the very real possibility that it might simply go nowhere or that it might not work as you expect?

I've struggled the most with this particular risk. The others I understand better how to mitigate; there are worst-case scenarios that can be dealt with. But this one is more existential, and it cuts across a broader swathe of professional responsibilities. If I start this project, and it doesn't work out, how can I justify the time I spent?

Even more to the point, if I start a new project, how can I justify spending the time that it needs to get bootstrapped, to reach a minimum viability, to even demonstrate some of its promise? Many of us have seen this before, and I confess I've even participated, for good and ill. In my case, I've seen this conflict arise around scientific platforms. Imagine a piece of software, used for knowledge production (and thus aligned with the economic and reputational interests of its users), that perhaps hasn't been updated with new design patterns, new languages, or modern libraries. Maybe it has started to buckle under design decisions made early in the project's career.

The opportunity to rebuild, redesign, and start afresh is enticing. But for me, at least, there's an undercurrent of dread: if I can't undertake this entire project by myself, how do I best set it up for success if I can't commit time to it? How do I build it out in such a way that others can participate, that it can show immediate results and be minimally disruptive?

Mitigating this type of risk is much more difficult because it requires both external and internal preparation. We have to design our projects in such a way as to minimize disruption to those who might be using them—chiefly, ourselves—which requires discipline in development, ensuring changes are as modular as possible and that each new state of the project is as working and functional as possible. If your project has an existing user base, it requires that disruptions are minimal and well-communicated, with stakeholders invested in the changes and the reasoning behind them.

How do we undertake something large, intimidating, overwhelming, knowing that it might be a failure down the road, or that it may never pay off?

To this, I don't have a ready, easy answer. For some reason, computational work seems to draw out different emotions than experimental work; experiments are by their nature “experimental,” sometimes without a known endpoint, and sometimes carrying with them the risk of the unknown result.

I can only share the strategies that work for me: faced with a large project, overwhelming and scary, focusing on the creative urge can be comforting. But this only lasts for a while, until infrastructure, maintenance, and integration are required to keep pushing the project forward to its conclusion. In these cases, I find myself struggling, and the most successful strategy I've found has been to foster collaboration. Building something can be isolating and lonely, but reducing that isolation can build social capital, share the burden of the project, and even guide things forward in unforeseen ways.

These risks and strategies are by no means the only ones we encounter in the process of scientific programming, and I certainly don't think that I've identified all the possible strategies for addressing them. I invite readers to share their thoughts on these and other approaches: When faced with difficulties or challenges, how have you addressed them? What strategies have you undertaken? What worked and what didn't? ■

---

**Matthew Turk** is a research scientist at NCSA at the University of Illinois. His research interests include data analysis, computational science, and scientific software communities. Turk received a PhD in physics from Stanford. Contact him at [matthewturk@gmail.com](mailto:matthewturk@gmail.com).

*This article originally appeared in  
Computing in Science & Engineering, vol. 18,  
no. 6, 2016.*

myCS

Read your subscriptions through  
the myCS publications portal at  
<http://mycs.computer.org>.



Containers provide something we've been trying to achieve for years: a standard application architecture that offers both managed distribution and service orientation.

Most compelling right now is containers' portability advantage. However, I suspect we'll discover more value over time. In fact, I suspect that containers will become a part of most IT shops, no matter whether they're moving to the cloud or not.<sup>2</sup>

### Defining a New Value for Containers

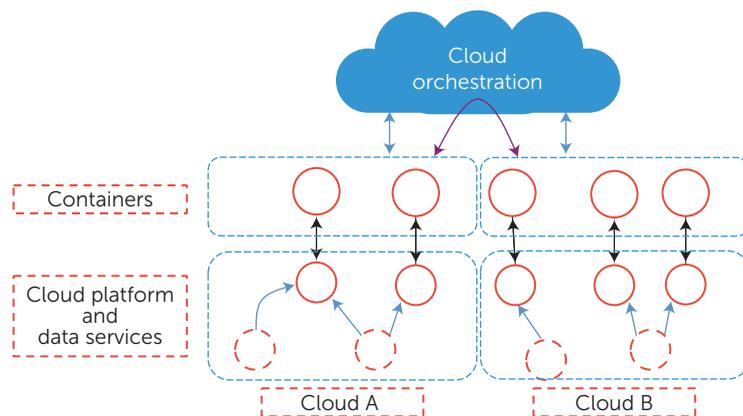
Containers are predicated on the goal of deploying and managing n-tier application designs. By their nature, containers manage n-tier application components, e.g., database servers, application servers, web servers, etc., at the operating system level. Indeed, portability is inherent because all operating system and application configuration dependencies are packaged and delivered inside a container to any other operating system platform. Containers are preferable to virtual machines here because they share compute platform resources very well whereas virtual machine platforms tend to acquire and hold resources on a machine-by-machine basis.<sup>3</sup>

In essence, containers can move from cloud to cloud and system to system, and thus can also provide automation for this process. In other words, we not only can leverage containers, but also can have them automatically "live migrate" from cloud to cloud as needed to support the application's requirements.

At the center of the container evolution is a cloud orchestration layer that can provision the infrastructure required to support the containers, as well as perform the live migration and monitor their health after the migration occurs (see Figure 1).

The concepts of autoprovisioning and automation are often promoted within modern cloud computing development but are elusive in practice. These concepts have a few basic features and advantages.

First is the ability to reduce complexity by leveraging container abstractions. Containers remove the dependencies on the underlying infrastructure ser-



**FIGURE 1.** A cloud orchestration layer oversees the infrastructure supporting live migration of containers.

vices, which reduces the complexity of dealing with those platforms. Containers are truly small platforms that support an application or an application's services that sit inside of a well-defined domain.

The second advantage is the ability to leverage automation with containers to maximize their portability, and thus their value. Through the use of automation, we script things we could also do manually, such as migrating containers from one cloud to another. We can also reconfigure communications between containers, such as tiered services, or data service access. However, today it's much harder to guarantee portability and application behavior when using automation. Indeed, automation often relies on many external dependencies that can break at any time, and thus remains a problem. However, it's indeed solvable.

Another advantage is the ability to provide better security and governance services by placing those services around, rather than within, containers. In many instances, security and governance services are platform-specific, not application-specific. Placing security and governance services outside of the application domain provides better portability and less complexity during implementation and operations.

Better distributed computing capabilities can also be provided since an application can be divided into many different domains, all residing with containers. These containers can be run on any number

of cloud platforms, including those that provide the most cost and performance efficiencies, and therefore applications can be distributed and optimized as to their use of the platform from within the container. For example, an I/O-intensive portion of the application could run on a bare metal cloud that provides the best performance, while a compute-intensive portion of the application runs on a public cloud that provides the proper scaling and load balancing. Perhaps even a portion of the application could run on traditional hardware and software. They all work together to form the application, and the application is separated into components that can be optimized.

Finally, there's the ability to provide automation services that offer policy-based optimization and self-configuration. None of this works without providing an automation layer that can "automagically" find the best place to run the container, as well as

noted in an email, portability of container clustering and orchestration is likely to quickly become the bottleneck.

### Making the Business Case

The problem with technical assertions is that they need to define a business benefit to be accepted by the industry as a best practice. The technical benefits I've defined need to be translated into direct business benefits that provide a quick return on investment.

One business benefit is the ability to automatically find least-cost cloud providers. Part of the benefit of moving from cloud to cloud is that you can leverage this portability to find the least-cost provider. Assuming most things are equal, the applications within a set of containers can live migrate to a cloud that offers price advantages for similar types of cloud services, such as storage.

For example, an inventory control application that exists within a dozen or so containers might have some storage-intensive components that cost \$100,000 a month on AWS. However, Google charges \$50,000 a month for the same types of resources. Understanding this configuration possibility within the orchestration layer, the containers can automigrate/live migrate to the new cloud where there's a 50 percent savings. If Google raises its prices and AWS lowers theirs, the reverse could occur.

The problem with technical assertions is that they need to define a business benefit to be accepted by the industry as a best practice.

deal with the changes in the configurations, and other things specific to the cloud platforms where the containers reside.

However, we've learned that n-tier applications have inherent limitations. "They are designed to scale up with very little focus paid on scaling down and no attention paid to scaling out or in. They typically are rife with single points of failure and tend to manage their own state via the use of cluster-style computing. Each tier of the n-tiered architecture must be scaled independently of the other tiers."<sup>3</sup>

Also, keep in mind that the automation/orchestration required will not always be portable. Indeed, that's likely the new lock-in layer; once you've built out the operational side, how easy is it to migrate from cloud to cloud? As Lori MacVittie of F5.com

These automation concepts also support better reliability. We've all done business cases around up-time and down-time. In some instances, businesses can lose as much as \$1 million an hour when systems aren't operating. Even if the performance issue lasts for only an hour or two, the lost productivity can move costs well into thousands of dollars per minute.

This architecture shown in Figure 1 can help avoid outages and related performance issues by opening other cloud platforms where the container workloads can relocate if issues occur on the primary clouds. For example, if AWS suffers an outage, the containers can be relocated to Google in a matter of minutes, where they can operate once again until the problem is resolved. You might

choose to run redundant versions of the containers on both clouds, supporting an active/active type of recovery platform.

## Facing Realities

Containers might sound like distributed application nirvana. They certainly offer a better way to utilize emerging cloud-based platforms. However, there are many roadblocks in front of us and a lot of work to be done.

We need to consider the fact that current technology can't provide this type of automation. Although it can certainly manage machine instances, even containers, using basic policy and scripting approaches, automatically moving containers from cloud to cloud using policy-driven automation, including autoconfiguration and autolocalization, isn't there yet.

Also, we've only just begun our Docker container journey. We still have a lot to learn about the technology's potential as well as its limitations. As we learned from the use of containers and distributed objects from years ago, the only way this technology can provide value is through coordinating clouds that support containers. Although having a standard here is great, history shows that vendors and providers tend to march off in their own proprietary directions for the sake of market share. If that occurs, all is lost.

The final issue is complexity. It only seems like we're making things less complex. Over time, the use of containers as the means of platform abstraction will result in applications that morph toward architectures that are much more complex and distributed. Moving forward, it might not be unusual to find applications that exist in hundreds of containers, running on dozens of different models and brands of cloud computing. The more complex these things become, the more vulnerable they are to operational issues.

## ALL THINGS CONSIDERED, CONTAINERS MIGHT BE A MUCH BETTER APPROACH TO BUILDING APPLICATIONS ON THE CLOUD.

PaaS and IaaS clouds will still provide the platform foundations and even development capabilities. However, these things will likely commoditize over

time, moving from a true platform to good container hosts. It will be interesting to see if the larger providers want to take on that role. Considering provider interest in Docker, that indeed could be their direction.

The core question now: if this is the destination of this technology and application hosting on cloud-based platforms, should I redirect resources toward this new vision? I suspect that most enterprises already have their hands full with the great cloud migration. However, as we get better at cloud application architectures using approaches that better account for both automation and portability, we'll eventually land on containers.

## Acknowledgments

Part of this article was derived from research I've done at Gigaom Pro, which is now out of business.

## References

1. Rightscale, *State of the Cloud Report*, 2016; <https://www.rightscale.com/lp/state-of-the-cloud?campaign=701700000015euX>.
2. D. Linthicum, "Fad? No, Containers Are Here to Stay," *InfoWorld*, 12 Feb. 2016; [www.infoworld.com/article/3032164/cloud-computing/fad-no-containers-are-here-to-stay.html](http://www.infoworld.com/article/3032164/cloud-computing/fad-no-containers-are-here-to-stay.html).
3. D. Linthicum, "Containers Are Designed for an Antiquated Application Architecture," *Container J.*, 5 June 2015; <http://containerjournal.com/2015/06/05/containers-are-designed-for-an-antiquated-application-architecture>.

---

**DAVID S. LINTHICUM** is senior vice president of Cloud Technology Partners. He also frequently writes for *InfoWorld* on deep technology subjects. His research interests include complex distributed systems, including cloud computing, data integration, service-oriented architecture, Internet of Things, and big data systems. Contact him at [david@davidlinthicum.com](mailto:david@davidlinthicum.com).

*This article originally appeared in  
IEEE Cloud Computing, vol. 3, no. 6, 2016.*

# Careers in Artificial Intelligence

**F**or this *ComputingEdge* issue, we asked Vanderbilt University professor of computer science and computer engineering Douglas Fisher about career opportunities in artificial intelligence. Fisher also serves as director for outreach, education, diversity, and synthesis for CompSustNet, a US National Science Foundation–sponsored research network used to explore new directions in computational sustainability. He authored the article “Recent Advances in AI for Computational Sustainability” in *IEEE Intelligent Systems*’ July/August 2016 issue.

**ComputingEdge:** What careers in artificial intelligence will see the most growth during the next several years?

**Fisher:** Many AI jobs will be broadly concerned with designing and implementing automation. They’ll be distributed across many sectors, including retail, transportation, and healthcare. However, for the most part, they won’t be about full automation because experience suggests this is fraught with problems. Rather, integrating humans and AI

to do tasks that were previously done by humans or AI alone is probably the future.

So, AI professionals will be working on teams with HCI (human-computer interaction) professionals, and it’s likely that expertise in this area will be a plus. The products they produce will include mobile devices, as well as intelligent personal assistants that work with, rather than frustrate, their owners.

Machine learning will also continue to be a dominant AI subfield for some time. And there is plenty of room to improve AI in games and in social and environmental simulations.

**ComputingEdge:** What would you tell college students to give them an advantage over the competition?

**Fisher:** Students should think about their interests and passions. Students who are in computer science for “guaranteed” jobs may get work but not necessarily a career they love. Well-meaning but, I think, misguided faculty members sometimes promote computer science for its job potential,

without considering other sources of motivation such as a love of technology or the possibility of advancing the social good.

Students should consider pairing their AI and computer-science interests with something else—perhaps a love of foreign languages and cultures, medicine and health, psychology, neuroscience, the law, journalism and writing, music, film, or other types of art. This can open students' eyes to what's possible in an AI career and make them more attractive to companies.

Internships are very important, too. These can include summer undergraduate research projects with faculty at their school or at another institution. Underclass students who don't feel ready for a full-blown computer-science internship might want to do something like teaching programming to kids at a summer camp. I think many companies appreciate a diversity of internship experience.

**ComputingEdge:** What should applicants keep in mind when applying for artificial-intelligence-related jobs?

**Fisher:** Understand the job, including its requirements, obligations, and freedoms. Is AI the only necessary area of expertise? Will they be trained in other skills on the job?

**ComputingEdge:** Name one critical mistake for young graduates to avoid when starting their careers?

**Fisher:** Don't forget about the future—both an individual's financial and family future, and a societal and environmental future that requires active contributions by people with brains, skills, and caring.

**ComputingEdge:** Do you have any learning experiences you could share that could benefit those just starting their careers?

**Fisher:** During the same semester as an undergraduate at the University of California, Irvine, I took a class in data structure techniques from

Thomas Standish in computer science and a class in cognition and memory from Tarow Indow in psychology. Studying how computers and humans represented and processed information led to small revelations that changed the direction of both my studies and my career. This goes to the point I made that combining computing and other fields can be a big win.

**C**omputingEdge's Lori Cameron interviewed Fisher for this article. Contact her at [l.cameron@computer.org](mailto:l.cameron@computer.org) if you would like to contribute to a future *ComputingEdge* article on computing careers. Contact Fisher at [douglas.h.fisher@vanderbilt.edu](mailto:douglas.h.fisher@vanderbilt.edu). 🍌

 Read your subscriptions through the myCS publications portal at <http://mycs.computer.org>.

## 2017 B. Ramakrishna Rau Award Call for Nominations

Honoring contributions to the computer microarchitecture field

**New Deadline: 1 May 2017**

 Established in memory of Dr. B. (Bob) Ramakrishna Rau, the award recognizes his distinguished career in promoting and expanding the use of innovative computer microarchitecture techniques, including his innovation in compiler technology, his leadership in academic and industrial computer architecture, and his extremely high personal and ethical standards.

**WHO IS ELIGIBLE?** The candidate will have made an outstanding innovative contribution or contributions to microarchitecture and use of novel microarchitectural techniques or compiler/architecture interfacing. It is hoped, but not required, that the winner will have also contributed to the computer microarchitecture community through teaching, mentoring, or community service.

**AWARD:** Certificate and a \$2,000 honorarium

**PRESENTATION:** Annually presented at the ACM/IEEE International Symposium on Microarchitecture

**NOMINATION SUBMISSION:** This award requires 3 endorsements. Nominations are being accepted electronically: [www.computer.org/web/awards/rau](http://www.computer.org/web/awards/rau)

**CONTACT US:** Send any award-related questions to [awards@computer.org](mailto:awards@computer.org).

[www.computer.org/web/awards/rau](http://www.computer.org/web/awards/rau)

## CAREER OPPORTUNITIES

**DATA MESSAGING MIDDLEWARE SPECIALIST**, Warren, MI, General Motors. Monitor, support, debug & troubleshoot GSSM (Global Sales, Services & Marketing), GPSC (Global Part Supply Chain) departments' app interfaces supporting vehicle mfg & aftersales. Support ICVM (Inter Continent Vehicle Movement) invoice processing & GWM (Global Warranty Management) warranty claim processing. Execute day to day operation tasks inclg anlysis of failed interfaces in ETL tools such as IBM Datastage, ICAN, SRE & JCAPS. Perform & manage day-to-day monitoring sys, implementing responsive actions, & reporting on the IBM Information Bus messaging sys. Deploy interfaces on IBM Datastage Information Server 11.3. Renew certificates on EAI operated & owned servers. Master, Computer Science, Computer Networking or related. 6 mos exp as Computer Programmer, Information Technology Analyst or related, executing day to day operation tasks inclg anlysis of failed interfaces in ETL tools such as IBM Datastage, ICAN, SRE & JCAPS. Mail resume to Alicia Scott-Wears, GM Global Mobility, 300 Renaissance Center,

MC:482-C32-D44, Detroit, MI 48265, Ref#35904.

**HUMAN RESOURCES DATA STRATEGIST**, General Motors, Detroit, MI. Desg, dvlp & execute apps architectures (using PeopleSoft & Oracle Developer), policies, practices & procedures to properly manage the full lifecycle for GM People's Data (employee, contractor, suppliers, & person of interest) over main sys of record of all GM ops. Strategize & ensure implementation of global HR data qlty mgmt rules' definitions, discovery processes, remediation policies, & standards inclg data dictionary with global HR, Finance & IT, ensuring one global data definition regardless of sys of record. Establish global Oracle PeopleSoft Data Governance, MDM (Master Data Mgmt) & DQM (Data Qlty Monitoring) processes to ensure qlty, integrity & security of people data in HR systems. Establish standards & procedures for people data life cycle mgmt CRUD standards to ensure data is recorded, stored, extracted & retired from HR systems in a consistent manner.

Transform raw data into insightful information for critical decision making using Tableau & MS Power BI. Run data qlty scans & conduct data profiling using IBM Cognos, IBM InphoSphere Information Servers, Information Analyzer & Atacama Data Qlty analyzer tools, to identify outlier patterns or anomalies in systems data & determine structured solutions for unstructured fuzzy data problems. Bachelor, Computer Engrg, Systems Engrg, Information Technology or related. 12 mos exp as HR Data Strategist, HR Data Scientist or related, executing apps architectures using PeopleSoft, & procedures to manage full lifecycle People's Data (employee, contractor, suppliers, & person of interest) over main sys of record for a region, establishing MDM & DQM processes, & running data qlty scans & conducting data profiling using IBM Cognos to identify outlier patterns or anomalies in systems data & determine structured solutions for unstructured fuzzy data problems. Mail resume to Ref#46806, GM Global Mobility, 300 Renaissance Center, MC:482-C32-D44, Detroit, MI 48265.

### TECHNICAL

#### Oracle America, Inc.

has openings for

## TECHNICAL ANALYST-SUPPORT

positions in **Orlando, FL**.

Job duties include: Deliver post-sales support and solutions to the Oracle customer base while serving as an advocate for customer needs.

Apply by e-mailing resume to [mine.odok@oracle.com](mailto:mine.odok@oracle.com), referencing 385.19791.

Oracle supports workforce diversity.

### SOFTWARE

#### Oracle America, Inc.

has openings for

## SOFTWARE DEVELOPER

positions in **Morrisville, NC**.

Job duties include: Design, develop, troubleshoot and/or test/QA software.

Apply by e-mailing resume to [christopher.moore@oracle.com](mailto:christopher.moore@oracle.com), referencing 385.16880.

Oracle supports workforce diversity.

### TECHNICAL

#### Oracle America, Inc.

has openings for

## TECHNICAL ANALYST

positions in **Orlando, FL**.

Job duties include: Analyze user requirements to develop, implement, and/or support Oracle's global infrastructure.

Apply by e-mailing resume to [sharan.alva@oracle.com](mailto:sharan.alva@oracle.com), referencing 385.19760.

Oracle supports workforce diversity.

**IT PROJECT MGRS & SOFTWARE DEVS SOUGHT.** IT Proj Mgrs req MS or equiv in CS, Comp Apps, Engg or rlted & 12 mos' relevant indus exp. Will also accept BS or equiv in above fields & 5 yrs' progressively responsible relevant indus exp. S/ware Devs req BS or equiv in above fields & 2yrs' relevant indus exp. All positions based out of San Jose, CA HQ & subj to reloc to various unanticipated US locations. Mail resumes to: HR Mgr, Happiest Minds Technologies Pvt. Ltd., 2051 Junction Ave, Ste 208, San Jose, CA 95131.

**INTERSHIP DEVELOPER,** Brightstar US, Inc. (Libertyville, IL) to assist in the dvlpmt of Brightstar's next generation e-Commerce processes & srvcs for customers. Align to dsgn reqmts of the Brightstar systems (applicable) to clients are managed & successful for applics locally, regionally or globally dvlpd. Identify techn'l issues & provide resolution steps. Dsgn & dvlp systems using Intershop platform. Reqmts: Bachelor's deg, or foreign equiv, in Comp Sci or rlted field. Min. 5 yrs exp in job or overall s/ware dvlpmt exp w/ a min. 3 yrs working w/ Intershop Enfinity 6.x or Intershop

7. Solid understanding of SQL, Oracle SWL preferred. Exp w/ project implmnt using technologies such as Java, JSP, Servlet, JDBC, XML/XSLT, & HTML. Multiple positions avail. Resumes to: Sapna Rao, HR Manager, 850 Technology Way, Libertyville, IL 60048.

**PROGRAMMER ANALYST:** Analyze, Develop & Design User Interface using knowledge in Oracle Apps(11i & R12) AP, AR, GL, FA, CM modules, Oracle Developer 9i/10g/11g, OAF, Forms10g, Reports10g, XML publisher, PL/SQL developer, SQL\*Loader, UNIX, Java, TOAD, WF Builder ,video/audio streaming API's, MarkView, Advisor, Informatica, Visual Basic, ASP .Net, C#.Net, SQL Server & Business Objects. Must be willing to travel & relocate to unanticipated client locations throughout US. Reqs MS in comp sci, eng or rel. Mail resumes to Nitya Software Solutions Inc. 3100 Mowry Avenue, Suite 205, Fremont, CA 94538.

**LEAD INTEGRATION DEVELOPER F/T** (Poughkeepsie, NY) Position involves travel to various unanticipated worksites

up to 100% of the time anywhere in the United States. Must have Bach deg or the foreign equiv in Comp Sci, or related w/2 yrs of exp Designing, Building & Supporting integration interfaces that meets business requirements using WebSphere Process Server, WebSphere Application Server, WebSphere Portal Server v8.x, IBM Rational Software, BlueMix Cloud Platform, DevOps Tools, IBM Sterling Platform, Web Services, SOAP, Databases (IBM DB2 8.x,9.x), scripting skills in JavaScript. Basic knowledge of Java J2EE (JSF, Struts, Hibernate, EJB, JDBC, REST, JSP, JMS, JUnit) and different communication protocols (TCP/IP, HTTP, FTP, SMTP). Send resume: Indotronix Int'l Corp, Recruiting (AE), 331 Main St, Poughkeepsie, NY 12601.

**SYSTEM ANALYST:** BA in comp science + 6 mo. exp in web dev., online-store, mngmt system, CRM req'd. Design syst for inventory, price, orders, dev. online store mngmt syst, config new softw, implement quality standards, analyze data, design procedures. Mail CV to UF Corporation, Inc. 1601 S Main Street, High Point, NC 27260.

**QA ANALYST**

**Oracle America, Inc.**

has openings for

**QA ANALYST**

positions in **Redwood Shores, CA**

Job duties include: Responsible for developing, applying and maintaining quality standards for company products with adherence to both internal and external standards.

Apply by e-mailing resume to [steve.howard@oracle.com](mailto:steve.howard@oracle.com), referencing 385.19855.

Oracle supports workforce diversity.

**SOFTWARE**

**Oracle America, Inc.**

has openings for

**SOFTWARE DEVELOPER**

positions in **Bedford, MA.**

Job duties include: Design, develop, troubleshoot and/or test/QA software.

Apply by e-mailing resume to [Jeffrey.cioli@oracle.com](mailto:Jeffrey.cioli@oracle.com), referencing 385.19772.

Oracle supports workforce diversity.



**CLOUDERA, INC. IS RECRUITING FOR OUR PALO ALTO, CA OFFICE:**

**SOFTWARE ENGINEER #37532:** Design, develop, support & maintain software components in Cloudera's product offering.

**CUSTOMER OPERATION ENGINEER #37441:** Investigate product related issues both for individual customers & for common trends that may arise

**MAIL RESUME WITH JOB CODE # TO CLOUDERA ATTN: HR 1001 PAGE MILL RD., BLDG. 3 PALO ALTO, CA 94304**

## CAREER OPPORTUNITIES

**SENIOR QUALITY ASSURANCE ANALYST**, General Motors, Detroit, MI. Test, validate, &perform soft updates Over the Air on General Motors vehicle internal module such as OnStar core module, &external modules such as the center stack module &the engine control unit through soft platform. Perform functional testing, back-end testing, installation testing &sys integration testing. Train Qlty Assurance Analysts &Jr Analysts on vehicle domains, soft updates reqmts, &Vehicle soft platform. Document &recommend test process, design vehicle module test artifacts, &functional flows. Dvlp test plans, test reports, &qlty metrics. Design vehicle module test artifacts such as test plan, test report &qlty metrics. Maintain manual &automation test suites &update test scripts for changes &enhancements. Mentor team members on performing test activities. Execute web service testing using SOAP user interface tool, &automation testing using Selenium WebDriver 2.0, TestNG, JUnit, &Jenkins tools. Perform database validation by writing SQL queries &PL/SQL. Dvlp

Shell Scripts for UNIX platform to deploy &monitor application. Master, Electrical, Electronics or, Computer Engrg or related. 36 mos exp as Quality Consultant, Software Engineer, Quality Assurance Analyst or related, performing soft test estimation, test planning, test designing, test optimization &test execution. Will accept bachelor's or foreign equiv degree, in Electrical, Electronic, Computer Engrg or related, followed by at least 5 yrs of progressive exp in the specialty, in lieu of required education &exp. Will also accept any equally suitable combination of education, training, &/or exp which would qualify applicant to perform job offered. Mail resume to Ref#35608, GM Global Mobility, 300 Renaissance Center, MC:482-C32-D44, Detroit, MI 48265.

**SENIOR SOFTWARE ENGINEER-CONNECTED CAR**, General Motors, Detroit, MI. Gather reqmts, perform, &execute full lifecycle Vehicle Service Mgmt System &Over the Air (OTA) reflash applications dvlpmt inclgd vehicle modules

&feature functionality on passenger vehicles, using Java/J2EE, Oracle, Mysql, AngularJS, &Ajax. Perform software, code &reqmts analysis, software review, identification of code metrics, system risk, software reliability &OOPS analysis, &aspect oriented programming. Mentor younger dvlpmt teams, &review software being dvlpd to ensure the highest standards in code qlty, performance, security, &overall delivery to reqmts. Collaborate with Software Architects to contribute to multilayered software solution architecture &drive design framework alignment. Dvlp, design &unit test software applications. Dvlp, design &unit test software applications. Master, Information Technology, Computer Science, or related. 6 mos exp as Software Developer, Computer Programmer, or related, gathering reqmts &executing &supporting full lifecycle vehicle diagnostic data collection sys &OTA reflash applications dvlpmt, using Java, Oracle &Ajax. Mail resume to Ref#2875, GM Global Mobility, 300 Renaissance Center, MC:482-C32-D44, Detroit, MI 48265.

### SOFTWARE

## Oracle America, Inc.

has openings for

# SOFTWARE DEVELOPERS

positions in **Broomfield, Colorado.**

Job duties include: As a member of the software engineering division, assist in defining and developing software for tasks associated with the developing, debugging or designing of software applications or operating systems.

Apply by e-mailing resume to  
grant.burkhart@oracle.com,  
referencing 385.19628.

Oracle supports workforce diversity.

### TECHNOLOGY

## Oracle America, Inc.

has openings for

# SYSTEMS ANALYST-SUPPORT

positions in **Redwood Shores, CA.**

Job duties include: Analyze user requirements, procedures, and problems to automate or improve existing systems and review computer system capabilities, workflow, and scheduling limitations. Some travel may be required to work on projects at various, unanticipated sites throughout the United States. May telecommute from home.

Apply by e-mailing resume to  
sj.kitterman@oracle.com,  
referencing 385.16903.

Oracle supports workforce diversity.

### TECHNOLOGY

## Oracle America, Inc.

has openings for

# CONSULTING PROJECT DIRECTOR

positions in **Redwood Shores, CA.**

Job duties include: Plan, initiate, and manage information technology (IT) projects. Lead and guide the work of technical staff. Travel to various, unanticipated sites throughout the U.S. required. May telecommute from home.

Apply by e-mailing resume to  
annette.lawrence@oracle.com,  
referencing 385.7516. Oracle supports  
workforce diversity.

## Cisco Systems, Inc. is accepting resumes for the following positions:

**ALPHARETTA, GA: Network Engineer (Ref.# ALP3):** Responsible for the operational support of internal network systems.

**AUSTIN, TX: Customer Support Engineer (Ref.# AUS5):** Responsible for providing technical support regarding the company's proprietary systems and software.

**BELLEVUE/SEATTLE, WA: Product Manager (Ref.# SEA384):** Develop the internal business case and customer value proposition for new offers and changes to existing offers, based on an understanding of the market conditions, customer decision drivers, and competitive dynamics.

**BEAVERTON, OR: Software Engineer (Ref.# BEA1):** Responsible for the definition, design, development, test, debugging, release, enhancement or maintenance of networking software.

**BOXBOROUGH, MA: Technical Lead/Leader (Ref.# BOX3):** Lead engineering groups on projects to design, develop or test hardware or software products. **Technical Lead/Leader (Ref.# BOX23):** Lead engineering groups on projects to design, develop or test hardware or software products. Telecommuting permitted. **Software Engineer (Ref.# BOX1):** Responsible for the definition, design, development, test, debugging, release, enhancement or maintenance of networking software.

**BOSTON, MA: Product Manager, Engineering (Ref.# BOS4):** Responsible for managing the development and implementation of new product introduction engineering activities to meet production launch schedules, quality and cost objectives. Telecommuting permitted.

**DENVER, CO: Software/QA Engineer (Ref.# DEN2):** Debug software products through the use of systematic tests to develop, apply, and maintain quality standards for company products. **Software Engineer (Ref.# DEN7):** Responsible for the definition, design, development, test, debugging, release, enhancement or maintenance of networking software. **Network Consulting Engineer (Ref.# DEN8):** Responsible for the support and delivery of Advanced Services to company's major accounts.

**HOUSTON, TX: Network Consulting Engineer (Ref.# HOU4):** Responsible

for the support and delivery of Advanced Services to company's major accounts.

**IRVINE, CA: Consulting Systems Engineer (Ref.# IRV15):** Provide specific end-to-end solutions and architecture consulting, technical and sales support for major account opportunities at the theater, area, or operation level. Telecommuting Permitted.

**JACKSONVILLE, FL: Solutions Architect (Ref.# JAC2):** Responsible for IT advisory and technical consulting services development and delivery. Travel may be required to various unanticipated locations throughout the United States.

**LAWRENCEVILLE, GA: Network Consulting Engineer (Ref.# LV16):** Responsible for the support and delivery of Advanced Services to company's major accounts. Telecommuting permitted.

**RESEARCH TRIANGLE PARK, NC: Software Engineer (Ref.# RTP3):** Responsible for the definition, design, development, test, debugging, release, enhancement or maintenance of networking software. **Software Engineer (Ref.# RTP773):** Responsible for the definition, design, development, test, debugging, release, enhancement or maintenance of networking software. Telecommuting permitted. **Customer Support Engineer (Ref.# RTP1):** Responsible for providing technical support regarding the company's proprietary systems and software.

**RICHARDSON, TX: Software Engineer (Ref.# RIC3):** Responsible for the definition, design, development, test, debugging, release, enhancement or maintenance of networking software.

**SAN DIEGO, CA: User Experience Specialist (Ref.# SD4):** Identify user interaction requirements and develop user experience interface specifications and guidelines.

**SAN FRANCISCO, CA: CNG Member of Technical Staff (Ref.# SF9):** Design, implement, and test software for a web application used by our customers for IT management.

**SAN JOSE/MILPITAS/SANTA CLARA, CA: Customer Support Engineer (Ref.# SJ3):** Responsible for providing technical support regarding the company's proprietary systems and software. **Test Engineer (Ref.# SJ16):** Build test equipment and test diagnostics for new products based on

manufacturing designs. **Technical Marketing Engineer (Ref.# SJ15):** Responsible for enlarging company's market and increasing revenue by marketing, supporting, and promoting company's technology to customers. **IT Engineer (Ref.# SJ7):** Responsible for development, support and implementation of major system functionality of company's proprietary networking products. **Manager, Engineering (Ref.# SJ4):** Schedule and evaluate the resources required for multiple projects in terms of human resources and hardware equipment allocation. **Network Consulting Engineer (Ref.# SJ107):** Responsible for the support and delivery of Advanced Services to company's major accounts. Travel may be required to various unanticipated locations throughout the United States. **Technical Marketing Engineer (Ref.# SJ178):** Responsible for enlarging company's market and increasing revenue by marketing, supporting, and promoting company's technology to customers. Travel may be required to various unanticipated locations throughout the United States. **Network Engineer (Ref.# SJ57):** Responsible for the operational support of internal network systems. **Network Consulting Engineer (Ref.# SJ9):** Responsible for the support and delivery of Advanced Services to company's major accounts. **Data Engineer (Ref.# SJ829):** Implement complex big data projects with large sets of data collected from various endpoints in the network and turn this information into useful insights. **Corporate Development Engineer (Ref.# SJ368):** Configure and troubleshoot routers and networking equipment. **Network Consulting Engineer (Ref.# SJ525):** Responsible for the support and delivery of Advanced Services to company's major accounts. Telecommuting permitted.

**WALTHAM, MA: Software Engineer (Ref.# WAL2):** Responsible for the definition, design, development, test, debugging, release, enhancement or maintenance of networking software.

**PLEASE MAIL RESUMES WITH REFERENCE NUMBER TO CISCO SYSTEMS, INC., ATTN: G51G, 170 W. Tasman Drive, Mail Stop: SJC 5/1/4, San Jose, CA 95134. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.**

[www.cisco.com](http://www.cisco.com)



# Microsoft

Microsoft Corporation currently has the following openings (job opportunities available at all levels, e.g., Principal, Senior and Lead levels):

## ALPHARETTA, GA

**Consultant:** Deliver design, planning, and implementation services that provide IT solutions to customers and partners. Requires travel throughout U.S. up to 75% with work to be performed at various worksites throughout the U.S. <https://jobs-microsoft.icims.com/jobs/7181/go/job>

## TEMPE, AZ

**Cloud Solution Architects/Solution Architects:** Architect software, platform, services, hardware or technology solutions. [http://bit.ly/MSJobs\\_Tech\\_Solns](http://bit.ly/MSJobs_Tech_Solns)

## IRVINE, CA

**Premier Field Engineers:** Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Requires domestic travel up to 75% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. [http://bit.ly/MS-Jobs\\_Support\\_Delivery](http://bit.ly/MS-Jobs_Support_Delivery)

**Premier Field Engineer:** Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Requires travel up to 50% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7246/go/job>

**Technology Solutions Professional:** Improve the Enterprise Mobility business metrics (revenue and scorecard) through excellence in technical sales strategy and execution. Requires travel throughout U.S. up to 50% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7144/go/job>

## LOS ANGELES, CA

**Technical Account Managers:** Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. Telecommuting permitted. [http://bit.ly/MSJobs\\_Delivery\\_Relationship\\_Mgmt](http://bit.ly/MSJobs_Delivery_Relationship_Mgmt)

**Senior Consultant:** Deliver design, planning, and implementation services that provide IT solutions to customers and partners. Requires domestic and international travel up to 50%. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7139/go/job>

**Technical Evangelist:** Collaborate with sales teams to understand customer requirements, promote the sale of products, and provide sales support. Requires domestic

and international travel up to 25%. <https://jobs-microsoft.icims.com/jobs/6753/go/job>

## MOUNTAIN VIEW, PALO ALTO, SUNNYVALE, CA

**Applied Scientist:** Utilize knowledge in applied statistics and mathematics to handle large amounts of data using various tools. [http://bit.ly/MSJobs\\_Data\\_Applied\\_Science](http://bit.ly/MSJobs_Data_Applied_Science)

**Data Scientist:** Manipulate large volumes of data, create new and improved techniques and/or solutions for data collection, management and usage. [http://bit.ly/MSJobs\\_Data\\_Applied\\_Science](http://bit.ly/MSJobs_Data_Applied_Science)

**Design Verification/Validation Engineers:** Responsible for ensuring the quality of Microsoft hardware products. [http://bit.ly/MSJobs\\_Hardware\\_Design\\_Verification\\_Eng](http://bit.ly/MSJobs_Hardware_Design_Verification_Eng)

**Hardware Dev., Test or Design Engineers, Hardware Engineers, Electrical Engineers, Design Engineers (all levels, including Leads and Managers):** Design, implement and test computer hardware that add strategic value to the company. ([http://bit.ly/MSJobs\\_Hardware\\_Dev\\_Eng](http://bit.ly/MSJobs_Hardware_Dev_Eng))([http://bit.ly/MSJobs\\_Electrical\\_Eng](http://bit.ly/MSJobs_Electrical_Eng))

**Hardware Dev., Test or Design Engineers, Hardware Engineers, Electrical Engineers, Design Engineers (all levels, including Leads and Managers):** Design, implement and test computer hardware. Requires Domestic and International travel up to 25%. ([http://bit.ly/MSJobs\\_Hardware\\_Dev\\_Eng](http://bit.ly/MSJobs_Hardware_Dev_Eng)) ([http://bit.ly/MSJobs\\_Electrical\\_Eng](http://bit.ly/MSJobs_Electrical_Eng))

**Solutions Sales Professional/Specialist/Technology Solutions Professionals:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. [http://bit.ly/MSJobs\\_Solution\\_Sales](http://bit.ly/MSJobs_Solution_Sales)

**Solutions Specialist:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7383/go/job>

**Sr. Quality and Reliability Engineer – Packaging:** Design, implement and test computer hardware products that add strategic value to the company. <https://jobs-microsoft.icims.com/jobs/7704/go/job>

## CAMBRIDGE, MA

**Senior Technical Evangelist:** Collaborate with sales teams to understand customer requirements, promote the sale of products, and provide sales support. Requires domes-

tic travel up to 25%. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7131/go/job>

**Solutions Sales Specialist:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Requires domestic travel up to 25%. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7927/go/job>

## SAN FRANCISCO, CA

**Data Scientist:** Manipulate large volumes of data, create new and improved techniques and/or solutions for data collection, management and usage. [http://bit.ly/MSJobs\\_Data\\_Applied\\_Science](http://bit.ly/MSJobs_Data_Applied_Science)

**Designers:** Develop user interface and user interaction designs, prototypes and/or concepts for business productivity, entertainment or other software or hardware applications. [http://bit.ly/MSJobs\\_Design](http://bit.ly/MSJobs_Design)

**Premier Field Engineers:** Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Requires domestic travel up to 75% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. [http://bit.ly/MS-Jobs\\_Support\\_Delivery](http://bit.ly/MS-Jobs_Support_Delivery)

**Solutions Sales Professional/Specialist/Technology Solutions Professionals:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. [http://bit.ly/MSJobs\\_Solution\\_Sales](http://bit.ly/MSJobs_Solution_Sales)

**Solution Specialist:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Requires domestic and international travel up to 25%. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7708/go/job>

**Solutions Specialist:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Requires domestic travel up to 25%. <https://jobs-microsoft.icims.com/jobs/7672/go/job>

## GUAYNABO, PR

**Premier Field Engineer, Senior:** Provide technical support to enterprise customers,

partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Travel up to 50% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7201/go/job>

#### FT. LAUDERDALE, FL

**Business Managers and Business Development Managers/Business Development and Strategy Analyst/Manager:** Develop business opportunities for sales of software and services. [http://bit.ly/MSJobs\\_Business\\_Development](http://bit.ly/MSJobs_Business_Development)

**Solutions Sales Professional/Specialist/Technology Solutions Professionals:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. [http://bit.ly/MSJobs\\_Solution\\_Sales](http://bit.ly/MSJobs_Solution_Sales)

**Business Operations Associate:** Analyze data to identify business insights related to business performance and growth. Requires international travel up to 25%. <https://jobs-microsoft.icims.com/jobs/7280/job>

**Delivery Manager:** Responsible for engaging with and managing accounts of customers utilizing Microsoft Dynamics AX and CRM business solution. Requires domestic and international travel up to 75%. <https://jobs-microsoft.icims.com/jobs/6867/go/job>

**Senior Technical Advisor:** Provide technical advice and support on issues experienced with Microsoft technologies. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7250/go/job>

**Solution Architect:** Architect software, platform, services, hardware or technology solutions. Requires international travel up to 25%. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/6815/job>

**Solution Specialist:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7398/go/job>

**Solutions Sales Specialist, Education:** Enhance Microsoft customer relationship from capability development perspective by articulating value of services & solutions & identifying competition gaps in targeted accounts. Requires travel up to 25% with work to be performed at various unknown worksites throughout the U.S. and Latin America. <https://jobs-microsoft.icims.com/jobs/6821/go/job>

#### CHICAGO, IL

**Cloud Solution Architects/Solution Architects:** Architect software, platform, services, hardware or technology solutions. Telecommuting permitted. [http://bit.ly/MSJobs\\_Tech\\_Solns](http://bit.ly/MSJobs_Tech_Solns)

**Cloud Solution Architect:** Responsible for architecting software, platform, services, hardware or technology solutions. Requires domestic and international travel up to 50%. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/6757/go/job>

**Solution Specialist:** Enhance Microsoft customer relationship from capability development perspective by articulating value of services and solutions and identifying competition gaps in targeted accounts. Requires domestic travel up to 25%. <https://jobs-microsoft.icims.com/jobs/6786/go/job>

#### DOWNERS GROVE, IL

**Associate Architect:** Architect software, platform, services, hardware, or technology solutions. Domestic travel required up to 50%. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/6776/go/job>

#### SOUTHFIELD, MI

**Solutions Sales Professional/Specialist/Technology Solutions Professionals:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. [http://bit.ly/MSJobs\\_Solution\\_Sales](http://bit.ly/MSJobs_Solution_Sales)

#### EDINA, MN

**Solutions Sales Professional/Specialist/Technology Solutions Professionals:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. [http://bit.ly/MSJobs\\_Solution\\_Sales](http://bit.ly/MSJobs_Solution_Sales)

#### CHARLOTTE, NC

**Support Engineers / Escalation Engineers:** Install, configure, support and troubleshoot issues related to Microsoft technologies. [http://bit.ly/MSJobs\\_Support\\_Eng](http://bit.ly/MSJobs_Support_Eng)

**Senior Technical Account Manager:** Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. Requires travel up to 25% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/6825/go/job>

**Technical Account Managers:** Assure productive use of Microsoft technologies, fo-

cus on delivery quality through planning and governance. [http://bit.ly/MSJobs\\_Delivery\\_Relationship\\_Mgmt](http://bit.ly/MSJobs_Delivery_Relationship_Mgmt)

#### FARGO, ND

**Support Engineers / Escalation Engineers:** Install, configure, support and troubleshoot issues related to Microsoft technologies. [http://bit.ly/MSJobs\\_Support\\_Eng](http://bit.ly/MSJobs_Support_Eng)

#### ISELIN, NJ

**Cloud Solution Architects/Solution Architects:** Architect software, platform, services, hardware or technology solutions. [http://bit.ly/MSJobs\\_Tech\\_Solns](http://bit.ly/MSJobs_Tech_Solns)

**Cloud Solution Architects/Solution Architects:** Architect software, platform, services, hardware or technology solutions. Telecommuting permitted. [http://bit.ly/MSJobs\\_Tech\\_Solns](http://bit.ly/MSJobs_Tech_Solns)

**Cloud Solution Architect:** Architect and deploy Microsoft cloud solutions for customers. Requires travel up to 25% with work to be performed at various unknown worksites throughout the U.S. <https://jobs-microsoft.icims.com/jobs/7164/go/job>

**Premier Field Engineer, SharePoint:** Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Requires travel up to 50% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7350/job>

**Solution Architect:** Engage with strategic customers to solve their business needs by conceiving, designing, and implementing digitally transformative technical architecture on Microsoft's public cloud offering (Azure). Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7242/go/job>

#### NEW YORK, NY

**Research Software Development Engineers (all levels):** Responsible for conducting applied research into new products and services through software engineering techniques. [http://bit.ly/MSJobs\\_Research\\_Software\\_Engineer](http://bit.ly/MSJobs_Research_Software_Engineer)

**Solutions Sales Professional/Specialist/Technology Solutions Professionals:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. [http://bit.ly/MSJobs\\_Solution\\_Sales](http://bit.ly/MSJobs_Solution_Sales)

CONTINUED ON NEXT PAGE...

## ...CONTINUED FROM PREVIOUS PAGE

**Account Technology Strategist:** Identify and analyze internal client and partner business needs, and translate needs into business requirements and value-added solutions and solution roadmaps. Requires travel up to 25% with work to be performed at various worksites throughout the U.S. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/6766/job>

**Associate Architect:** Architect solutions by mapping common customer business problems to reusable end-to-end technology solutions. Requires domestic and international travel up to 50%. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7170/go/job>

**Cloud Solution Architect:** Architect and deploy Microsoft cloud solutions for customers. Requires travel up to 25% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/6802/job>

**Principal Solution Specialist:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Requires domestic travel up to 25% to various unknown locations. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7265/go/job>

**Solution Sales Professional:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Requires travel up to 25% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7243/go/job>

**Solution Specialist:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Requires travel up to 25% to perform work at various unknown customer sites. <https://jobs-microsoft.icims.com/jobs/7380/job>

**Solutions Specialist:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Requires travel up to 50% with work to be performed at various unanticipated worksites throughout the U.S. Telecommut-

ing permitted. <https://jobs-microsoft.icims.com/jobs/6750/go/job>

**Technical Evangelist:** Collaborate with sales teams to understand customer requirements, promote the sale of products, and provide sales support. Requires domestic travel up to 25%. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7746/go/job>

**Technology Solutions Professional Devices:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7065/job>

### WILSONVILLE, OR

**Hardware Dev., Test or Design Engineers, Hardware Engineers, Electrical Engineers, Design Engineers (all levels, including Leads and Managers):** Design, implement and test computer hardware that add strategic value to the company. ([http://bit.ly/MSJobs\\_Hardware\\_Dev\\_Eng](http://bit.ly/MSJobs_Hardware_Dev_Eng))([http://bit.ly/MSJobs\\_Electrical\\_Eng](http://bit.ly/MSJobs_Electrical_Eng))

**Hardware Dev., Test or Design Engineers, Hardware Engineers, Electrical Engineers, Design Engineers (all levels, including Leads and Managers):** Design, implement and test computer hardware. Requires Domestic and International travel up to 25%. ([http://bit.ly/MSJobs\\_Hardware\\_Dev\\_Eng](http://bit.ly/MSJobs_Hardware_Dev_Eng)) ([http://bit.ly/MSJobs\\_Electrical\\_Eng](http://bit.ly/MSJobs_Electrical_Eng))

### HOUSTON, TX

**Technical Account Managers:** Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. Roving Employee - requires travel up to 100% with work to be performed at various unknown worksites throughout the U.S. [http://bit.ly/MSJobs\\_Delivery\\_Relationship\\_Mgmt](http://bit.ly/MSJobs_Delivery_Relationship_Mgmt)

### AUSTIN, TX

**Partner Business Evangelist:** Research market conditions in regional area and gather information to determine sales of Windows Azure. Requires Domestic and International travel up to 25%. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7530/job>

### IRVING, TX

**Support Engineers / Escalation Engineers:** Install, configure, support and troubleshoot issues related to Microsoft technologies. [http://bit.ly/MSJobs\\_Support\\_Eng](http://bit.ly/MSJobs_Support_Eng)

**Premier Field Engineer:** Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Requires travel up to 75% with work to be performed at various unknown worksites throughout the U.S. Telecommuting per-

mitted. <https://jobs-microsoft.icims.com/jobs/7178/go/job>

### MCALLEN, TX

**Hardware Dev., Test or Design Engineers, Hardware Engineers, Electrical Engineers, Design Engineers (all levels, including Leads and Managers):** Design, implement and test computer hardware that add strategic value to the company. ([http://bit.ly/MSJobs\\_Hardware\\_Dev\\_Eng](http://bit.ly/MSJobs_Hardware_Dev_Eng))([http://bit.ly/MSJobs\\_Electrical\\_Eng](http://bit.ly/MSJobs_Electrical_Eng))

**Hardware Dev., Test or Design Engineers, Hardware Engineers, Electrical Engineers, Design Engineers (all levels, including Leads and Managers):** Design, implement and test computer hardware. Requires Domestic and International travel up to 25%. ([http://bit.ly/MSJobs\\_Hardware\\_Dev\\_Eng](http://bit.ly/MSJobs_Hardware_Dev_Eng)) ([http://bit.ly/MSJobs\\_Electrical\\_Eng](http://bit.ly/MSJobs_Electrical_Eng))

### RESTON, VA

**Technical Account Managers:** Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. [http://bit.ly/MSJobs\\_Delivery\\_Relationship\\_Mgmt](http://bit.ly/MSJobs_Delivery_Relationship_Mgmt)

### REDMOND, WA

**Applied Scientist:** Utilize knowledge in applied statistics and mathematics to handle large amounts of data using various tools. [http://bit.ly/MSJobs\\_Data\\_Applied\\_Science](http://bit.ly/MSJobs_Data_Applied_Science)

**Business Managers and Business Development Managers/Business Development and Strategy Analyst/Manager:** Develop business opportunities for sales of software and services. [http://bit.ly/MSJobs\\_Business\\_Development](http://bit.ly/MSJobs_Business_Development)

**Business Process Manager:** Responsible for the design, implementation, and release of programs or projects. [http://bit.ly/MSJobs\\_Fin\\_Plan\\_Analy\\_Contr](http://bit.ly/MSJobs_Fin_Plan_Analy_Contr)

**Cloud Solution Architects/Solution Architects:** Architect software, platform, services, hardware or technology solutions. [http://bit.ly/MSJobs\\_Tech\\_Solns](http://bit.ly/MSJobs_Tech_Solns)

**Consultants:** Deliver design, planning, and implementation services that provide IT solutions to customers and partners. Requires domestic and international travel up to 25%. [http://bit.ly/MSJobs\\_Technical\\_Delivery](http://bit.ly/MSJobs_Technical_Delivery)

**Consultants:** Deliver design, planning, and implementation services that provide IT solutions to customers and partners. Roving Employee—requires travel up to 100% with work to be performed at various unknown worksites throughout the U.S. [http://bit.ly/MSJobs\\_Technical\\_Delivery](http://bit.ly/MSJobs_Technical_Delivery)

**Consultants:** Deliver design, planning, and implementation services that provide IT solutions to customers and partners. Roving Employee—requires travel up to 100% with

work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. [http://bit.ly/MSJobs\\_Technical\\_Delivery](http://bit.ly/MSJobs_Technical_Delivery)

**Content Developer/Engineer:** (All levels, including Leads and Managers) Responsible for the design, development, deployment, vision, and business strategy for content creation, acquisition, production, editorial, and publishing activities at Microsoft. [http://bit.ly/MSJobs\\_Content\\_Publishing](http://bit.ly/MSJobs_Content_Publishing)

**Data Scientist:** Manipulate large volumes of data, create new and improved techniques and/or solutions for data collection, management and usage. [http://bit.ly/MSJobs\\_Data\\_Applied\\_Science](http://bit.ly/MSJobs_Data_Applied_Science)

**Design Researchers:** Develop user interface and user interaction designs, prototypes and/or concepts for business productivity, entertainment or other software or hardware applications. [http://bit.ly/MSJobs\\_Design\\_Research](http://bit.ly/MSJobs_Design_Research)

**Design Verification/Validation Engineers:** Responsible for ensuring the quality of Microsoft hardware products. [http://bit.ly/MSJobs\\_Hardware\\_Design\\_Verification\\_Eng](http://bit.ly/MSJobs_Hardware_Design_Verification_Eng)

**Designers:** Develop user interface and user interaction designs, prototypes and/or concepts for business productivity, entertainment or other software or hardware applications. [http://bit.ly/MSJobs\\_Design](http://bit.ly/MSJobs_Design)

**Evangelists/Technical Evangelists:** Collaborate with sales teams to understand customer requirements, promote the sale of products, and provide sales support. [http://bit.ly/MSJobs\\_Tech\\_Evangelist](http://bit.ly/MSJobs_Tech_Evangelist)

**Evangelists/Technical Evangelists:** Collaborate with sales teams to understand customer requirements, promote the sale of products, and provide sales support. Requires travel throughout the U.S. up to 25%. [http://bit.ly/MSJobs\\_Tech\\_Evangelist](http://bit.ly/MSJobs_Tech_Evangelist)

**Hardware Dev., Test or Design Engineers, Hardware Engineers, Electrical Engineers, Design Engineers (all levels, including Leads and Managers):** Design, implement and test computer hardware that add strategic value to the company. ([http://bit.ly/MSJobs\\_Hardware\\_Dev\\_Eng](http://bit.ly/MSJobs_Hardware_Dev_Eng))([http://bit.ly/MSJobs\\_Electrical\\_Eng](http://bit.ly/MSJobs_Electrical_Eng))

**Hardware Dev., Test or Design Engineers, Hardware Engineers, Electrical Engineers, Design Engineers (all levels, including Leads and Managers):** Design, implement and test computer hardware. Requires Domestic and International travel up to 25%. ([http://bit.ly/MSJobs\\_Hardware\\_Dev\\_Eng](http://bit.ly/MSJobs_Hardware_Dev_Eng)) ([http://bit.ly/MSJobs\\_Electrical\\_Eng](http://bit.ly/MSJobs_Electrical_Eng))

**Premier Field Engineers:** Provide technical support to enterprise customers, partners, internal staff or others on mission critical

issues experienced with Microsoft technologies. Roving Employee—requires travel up to 100% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. [http://bit.ly/MSJobs\\_Support\\_Delivery](http://bit.ly/MSJobs_Support_Delivery)

**Research Software Development Engineers (all levels):** Responsible for conducting applied research into new products and services through software engineering techniques. [http://bit.ly/MSJobs\\_Research\\_Software\\_Engineer](http://bit.ly/MSJobs_Research_Software_Engineer)

**Researchers/Scientists:** Conduct research and lead research collaborations that yield new insights, theories, analyses, data, algorithms, and prototypes and that advance state-of-the-art of computer science and engineering, as well as general scientific knowledge. [http://bit.ly/MSJobs\\_Research](http://bit.ly/MSJobs_Research)

**Service Engineers/Managers, Service Operations Engineers, and Systems/Operations Engineers/Site Reliability Engineer:** Research, design, develop, and test operating systems-level software, compilers, and network distribution software. ([http://bit.ly/MSJobs\\_Service\\_Engineering](http://bit.ly/MSJobs_Service_Engineering)) ([http://bit.ly/MSJobs\\_IT\\_Serv\\_Eng](http://bit.ly/MSJobs_IT_Serv_Eng)) ([http://bit.ly/MSJobs\\_IT\\_Serv\\_Ops](http://bit.ly/MSJobs_IT_Serv_Ops))

**Solutions Sales Professional/Specialist/Technology Solutions Professionals:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. [http://bit.ly/MSJobs\\_Solution\\_Sales](http://bit.ly/MSJobs_Solution_Sales)

**Support Engineers / Escalation Engineers:** Install, configure, support and troubleshoot issues related to Microsoft technologies. [http://bit.ly/MSJobs\\_Support\\_Eng](http://bit.ly/MSJobs_Support_Eng)

**Architect:** Deliver design, planning, and implementation services that provide IT solutions to customers and partners. <https://jobs-microsoft.icims.com/jobs/7389/job>

**Associate Architect:** Plan, develop and market solutions that enable Microsoft Sellers, architects, and consultants to sell and deliver our products and services. <https://jobs-microsoft.icims.com/jobs/7257/job>

**Associate Architect:** Design and develop business intelligence, big data and advanced analytics-based IT solutions for Microsoft customers. Requires travel up to 25% with work to be performed at various unknown worksites throughout the U.S. <https://jobs-microsoft.icims.com/jobs/6789/job>

**Associate Architect:** Architect and deploy complex solutions mapping customer business problems with end-to-end technology solutions. Requires domestic and international travel up to 50%. Telecommuting per-

mitted. <https://jobs-microsoft.icims.com/jobs/6819/job>

**Audience Engagement Development Engineer:** Responsible for building audience measurement and engagement tools to support the organization's effort in understanding audience experiences relative to the communications and marketing objectives. <https://jobs-microsoft.icims.com/jobs/7347/job>

**B2B Global Demand Center Analyst:** Responsible for converting high volumes of complex data into actionable insights. <https://jobs-microsoft.icims.com/jobs/6862/go/job>

**Business Analytics Specialist:** Provide insights, facilitate data-driven testing and learning, and develop ways to optimize Windows Devices Group (WDG) performance marketing campaigns. <https://jobs-microsoft.icims.com/jobs/7377/go/job>

**Business Analytics Specialist:** Identify strategic areas of focus for insights and analytics. <https://jobs-microsoft.icims.com/jobs/7137/go/job>

**Business Analytics Specialist:** Provide deep level data analytics and modeling for search advertising platform. <https://jobs-microsoft.icims.com/jobs/6806/job>

**Business Analytics Specialist, Digital Technology & Operations:** Analyze and optimize Business Intelligence systems and procedures. <https://jobs-microsoft.icims.com/jobs/7148/go/job>

**Business Development Analyst:** Develop and analyze business opportunities for sales of software and services. Requires domestic and international travel up to 25%. <https://jobs-microsoft.icims.com/jobs/7230/go/job>

**Business Manager:** Develop business opportunities for sales of software and services. Requires domestic and international travel up to 25%. <https://jobs-microsoft.icims.com/jobs/6758/go/job>

**Channel Executive OEM:** Build and maintain business relationships with Original Equipment Manufacturers (OEM) to assess needs, determine requirements, and deliver new Agreements and net new opportunity revenue. <https://jobs-microsoft.icims.com/jobs/6756/go/job>

CONTINUED ON NEXT PAGE...

## ...CONTINUED FROM PREVIOUS PAGE

**Channel Executive-Hosting:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Requires travel throughout U.S. up to 25% with work to be performed at various unknown worksites throughout the U.S. <https://jobs-microsoft.icims.com/jobs/6775/go/job>

**Cloud Solution Architect:** Architect and deploy Microsoft Azure cloud solutions for portfolio partners. <https://jobs-microsoft.icims.com/jobs/7684/go/job>

**Cloud Solution Architect:** Architect software, platform, services, hardware or technology solutions. Requires domestic and international travel up to 25%. Telecommuting Permitted. <https://jobs-microsoft.icims.com/jobs/7067/go/job>

**Cloud Solution Architect:** Design, plan, and implement cloud deployments to move customer workloads to Azure. Roving Employee—requires travel up to 100% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7184/go/job>

**Data & Applied Scientist; Apply scientific methodology and algorithms to large volumes of data using various tools.** <https://jobs-microsoft.icims.com/jobs/7688/go/job>

**Design Researcher II:** Develop user interface and user interaction designs, prototypes and/or concepts for business productivity, entertainment or other software or hardware applications. Requires domestic and international travel up to 25%. <https://jobs-microsoft.icims.com/jobs/6961/job>

**IT Audit Group Manager:** Manage a portfolio of IT audit projects including driving value to the business, quality, and execution excellence. Requires domestic travel up to 50%. <https://jobs-microsoft.icims.com/jobs/7365/go/job>

**Marketing Operations Automation Specialist:** Utilize modern big data management tools, marketing analytics methods, CRM tools, and marketing automation technology to prepare findings in marketing trends and assist the design of marketing plans. <https://jobs-microsoft.icims.com/jobs/7355/job>

**Optical Engineer II:** Design tools, technologies and processes for next generation optical systems. <https://jobs-microsoft.icims.com/jobs/7420/go/job>

**Packaging Engineer:** Develop, qualify and implement structural packaging for hard-

ware products that add strategic value to the company. <https://jobs-microsoft.icims.com/jobs/7046/job>

**Premier Field Engineer:** Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Requires travel throughout U.S. up to 50% with work to be performed at various unknown worksites throughout the U.S. <https://jobs-microsoft.icims.com/jobs/7211/go/job>

**Premier Field Engineer:** Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Requires international and domestic travel up to 50%. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7037/go/job>

**Premier Field Engineer:** Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Travel up to 50% with work to be performed at various unknown worksites throughout the US. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/6751/go/job>

**Premier Field Engineer:** Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Requires travel up to 25% with work to be performed at various unknown worksites throughout the U.S. <https://jobs-microsoft.icims.com/jobs/6752/go/job>

**Principal Researcher:** Conduct research and lead research collaborations that yield new insights, theories, analyses, data, algorithms and prototypes and that advance state-of-the-art of computer science and engineering, as well as general scientific knowledge. Requires international and domestic travel up to 25%. <https://jobs-microsoft.icims.com/jobs/7064/job>

**Product Intelligence Manager:** Build business process and business intelligence solutions based upon different telemetry systems. <https://jobs-microsoft.icims.com/jobs/6788/go/job>

**Senior Camera Module Engineer:** Specify, design, implement, and verify hardware for microcomputers. Requires domestic and international travel up to 25%. <https://jobs-microsoft.icims.com/jobs/6949/go/job>

**Senior Designer:** Support the creation and improvement of user experiences for new products, and mature all successful industrial design concepts through to product launch. <https://jobs-microsoft.icims.com/jobs/7582/go/job>

[jobs/7582/go/job](https://jobs-microsoft.icims.com/jobs/7582/go/job)

**Senior Industry Technology Strategist:** Work with commercial segments to help them successfully transform and deploy Microsoft technologies. <https://jobs-microsoft.icims.com/jobs/7251/go/job>

**Senior IT Service Operations:** Partner within Engineering organizations and other stakeholders to be continually improving the predictability of our service and enhancing our service offerings based on customer feedback balanced against business priorities. <https://jobs-microsoft.icims.com/jobs/7400/job>

**Senior Patent Engineer:** Perform analysis, reverse engineering, testing and investigations on software and information technology that are being deployed and used widely in the industry. <https://jobs-microsoft.icims.com/jobs/6773/job>

**Senior Premier Field Engineer:** Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Position requires travel up to 25% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/6891/go/job>

**Service Engineer 2:** Responsible for network designs, deployments, testing, and certifications; and operational escalation support of complex network infrastructure <https://jobs-microsoft.icims.com/jobs/7425/go/job>

**Solution Specialist:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Requires travel up to 25% throughout the U.S. <https://jobs-microsoft.icims.com/jobs/7168/go/job>

**Solutions Architect:** Engage with strategic customers to solve their business needs by conceiving, designing, and implementing transformative technical architecture on data and cloud-based platforms and related technologies. Roving Employee—requires travel up to 100% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7387/go/job>

**Solutions Specialist (Technology Solutions Professional):** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Requires travel throughout the U.S. up to 25%. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/6840/job>

**Support Engineer:** Install, configure, support and troubleshoot issues related to Microsoft technologies. Requires domestic travel up to 25% to perform work at various unknown customers' worksites. <https://jobs-microsoft.icims.com/jobs/6810/job>

**Support Planner:** Build assisted support experiences for customers inquiring about our product offerings or about an existing order. <https://jobs-microsoft.icims.com/jobs/7190/go/job>

**Technical Data Analyst:** Analyze, validate, and present data findings to answer high-impact business and product development questions. <https://jobs-microsoft.icims.com/jobs/6768/job>

**Technology Solutions Professional:** Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Requires domestic travel up to 25% to perform work at unknown customer sites. <https://jobs-microsoft.icims.com/jobs/7366/job>

**RENO, NV**

**Partners OPS Manager Launch:** Plan, initiate, and manage information technology (IT) projects. <https://jobs-microsoft.icims.com/jobs/7208/go/job>

**SACRAMENTO, CA**

**Premier Field Engineer, SharePoint Dev:** Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Requires local travel up to 25%. Telecommuting permitted. <https://jobs-microsoft.icims.com/jobs/7357/go/job>

**MADISON, WI**

**Research Software Development Engineers (all levels):** Responsible for conducting applied research into new products and services through software engineering techniques. [http://bit.ly/MS-Jobs\\_Research\\_Software\\_Engineer](http://bit.ly/MS-Jobs_Research_Software_Engineer)

**Multiple job openings are available for each of these categories. To view detailed job descriptions and minimum requirements, and to apply, visit the web-site address listed. EOE.**

**SOFTWARE ENGINEER,** Warren, MI, General Motors. Dvlp soft tools to support the dvlpmt of GM's in Over the Air (OTA) vehicle soft remote reflash dvlpmt kit for 2018-2020 &beyond model year passenger vehicle infotainment/radios &telematic systems, using C/C++, &performing root cause analysis on defects &provide fixes. Use network app protocols incldg HTTP, HTTPS, FTP, TCP/IP to troubleshoot, debug &fix OTA issues. Coordinate &participate in architecture, reqmt, design, code, &test case reviews. Dvlp soft using C/C++ &porting to Radio head units &Telematic Devices that run on QNX, Linux &Android operating systems. Use State Machine, Threads, Socket programming, Mutex, Semaphore &message queues for soft dvlpmt. Use SQLite for storing sys status over vehicle ignition cycles. Dvlp automation test scripts using perl for remote reflash unit test execution. Dvlp CANoE vector tool scripts to simulate Electronic Control Units. Use QNX Momentics as integrated dvlpmt environment. Use Vehicle Spy vector tool to monitor messages on CAN bus. Bachelor, Computer Science, Computer Engrg, or related. 60 mos exp as Software Engineer, Senior Software Engineer, Remote Reflash Engineer, using State Machine, Threads, Socket programming, Mutex, Semaphore &message queues for soft dvlpmt, &network app protocols incldg HTTP, HTTPS, FTP, TCP/IP to troubleshoot, debug &fix wireless or OTA issues. Exp must be post baccalaureate &progressive. Any suitable combination of education, training &exp is acceptable. Mail resume to Ref#2675, GM Global Mobility, 300 Renaissance Center, MC:482-C32-D44, Detroit, MI 48265.

**APPLICATION SUSTAIN ENGINEER,** General Motors, Detroit, MI. Deploy, operate, &enhance OnStar Infotainment Advanced Systems Dvlpmt Pilot Apps &operating environments such as Unix, Oracle, Java. Use &improve J2EE &Web Svcs-based app soft to analyze &perform OTA Reflash &data collection of OnStar Electronic Control Units (ECU) in vehicles communicating through Vehicle Computer Platform (VCP). Enhance current Pilot App product cmpts &lead dvlpmt efforts for new service offerings to troubleshoot pilot apps. Create certificates service request, generate Java keystore using openssl &keytool IUI for OnStar ADPA cmpts for vehicle communication using SMS &packet gateways, implementing certificates accordingly to VSMS environments &test them using VSMS. Use Jenkins &zuul configuration tools to build, deploy &operate the apps in N.A., Europe &China. Configure zabbix, kibana &log stash for log mgmt.

Create oracle &mysql database scripts to debug issues related to vehicle data collection &reflash. Create Perl script for app failover &maintain high availability of app. Create &request infrastructure team to provide CSS, firewall &IP tables rules for app to interact with vehicle &to external cmpt svces. Master, Electrical or Electronics Engineering or related. 36 mos exp as Network Infrastructure Engineer, Module Lead, Senior Software Engineer, Software Engineer or related, using &improving J2EE &Web Svcs-based app soft to analyze &perform OTA Reflash &data collection of OnStar ECU in vehicles communicating through VCP. Will accept bachelor's or foreign equiv degree, in Electrical or Electronics Engineering or related, followed by at least 5 yrs progressive exp in the specialty, in lieu of required education &exp. Will also accept any equally suitable combination of education, training, &/or exp which would qualify applicant to perform job offered. Mail resume to Ref#2783, GM Global Mobility, 300 Renaissance Center, MC:482-C32-D44, Detroit, MI 48265.

**STORAGE ENGINEER,** Warren, MI, General Motors Company. Perform operational support for 10PB storage allocation, resolving storage incidents, evaluating team member scripts, reporting on storage hardware, &performing regular capacity planning. Execute scripts for storage allocation on hardware such as EMC Vmax, Vplex, XtremIO, DMX-4/DMX-3, CX3-80 using symcli &Unisphere, create hyper mapping Luns, Lun masking, &configure RAID groups on Clariion LUN. Configure Zoning &troubleshoot storage &host related issues. Engr storage hardware for capacity to ensure storage provisioning, efficiency &maintenance. Desg SRDF/S (Symmetrix Remove Data Facility Synchronous) &Timefinder Clone/Snap on production sys for disaster recovery meeting RTO (Recovery time objectives) &RPO (Recovery Point Objectives) for SAP. Master, Computer Science, Information Systems, or related. 6 mos exp as Storage Administrator, Storage Engr or related, executing scripts for storage allocation on hardware such as EMC Vmax &Vplex using symcli &Unisphere, &creating Lun masking, or related. Mail resume to Ref#3493, GM Global Mobility, 300 Renaissance Center, MC:482-C32-D44, Detroit, MI 48265.

**VISITING FACULTY POSITION** in Computer Science (Non-Tenure Track) The Department of Computer Science at The College of New Jersey invites applications for a ten-month, non-tenure track visiting faculty position in computer science starting August 2017.

## CAREER OPPORTUNITIES

**SENIOR RESEARCHER**, Warren, MI, General Motors. Research, investigate, innovate, & dvlp new methods & techniques to automate (or complement techniques in) rigorous dvlpmt, verification, validation & testing of high-integrity embedded electronic control unit systems & soft that implement required automotive control features such as Automated Driving Sys Controllers, Automated steering sub sys, Power distribution sub sys, Vehicle controls, Interface gateways, Automatic Lane Changing & Obstacle Deduction & Avoidance features in autonomous vehicles for smoother & safer ride qlty, & conventional vehicle features such as Body Control Module, Engine Control Module, Airbag Control Module. Dvlp first-order predicate logic model to analyze safety impact of autonomous systems feature interfaces through concurrent sys predicate & propositional modelling, & anlys through Boolean satisfiability solver techniques. Perform verification & validation using Binary Decision Diagrams based on advanced formal verification methodology with

bounded constraint solvers & CONCOLIC testing, with traversal approaches combining concrete & symbolic simulation by means of local constraint solvers & Monte Carlo Simulation approaches. Dvlp techniques to check code conformance of auto generated code from parallel & hierarchical State Flow & Simulink models using the Model Based Testing (MBT) approaches combined with Automatic Test-case Generation (ATG) methods with enhanced model structural coverage. PhD, Computer Science, Information Technology, Computer Engrg, or related. 12 mos exp as Researcher, Scientist or Engineer, identifying & solving gaps in tools & methods in existing engrg anlys & dvlpg techniques, & methods for psgr vehicle features such as Automated steering sub sys, Power distribution sub sys, Automatic lane changing & warning & Obstacle deduction & collision avoidance in autonomous vehicles, dvlpg techniques to check code conformance of auto generated code from parallel & hierarchical State Flow & Simulink models using MBT

approaches combined with ATG methods with enhanced model structural coverage. Mail resume to Ref#2456, GM Global Mobility, 300 Renaissance Center, MC:482-C32-D44, Detroit, MI 48265.

**SENIOR USER EXPERIENCE & PRODUCT DESIGN.** Position available in New York, NY. Provide full stack services from information architecture, wire frames and flow charts, User Experience/User Interface and visual design. Lead internal and external teams on multi-platform projects in the design and development of best user experiences for large scale websites and applications. Conduct usability evaluations of existing and proposed designs in-person or using online platform [usertesting.com](http://usertesting.com). Create wireframes and design prototypes using Sketch, InVision, Adobe Creative Suite and Axure. Apply: L. Sawtelle, MIP F105, Massachusetts Mutual Life Insurance Company, 1295 State Street, Springfield, MA 01111; Please Reference Job ID: 708203400



Microsoft Corporation currently has the following openings (job opportunities available at all levels, e.g., Principal, Senior and Lead levels):

### IRVINE, CA

**Business Program Managers:** Plan, initiate, and manage technology and business projects. [http://bit.ly/MSJobs-Buss\\_Oper\\_Prog\\_Mgmt](http://bit.ly/MSJobs-Buss_Oper_Prog_Mgmt)

### MOUNTAIN VIEW, PALO ALTO, SUNNYVALE, CA

**Program Managers:** (All levels, including Leads and Managers) Coordinate program development of computer software applications, systems or services working with development and product planning teams. ([http://bit.ly/MSJobs\\_ProgMgr](http://bit.ly/MSJobs_ProgMgr)) ([http://bit.ly/MSJobs\\_HW\\_ProgMgr](http://bit.ly/MSJobs_HW_ProgMgr)) ([http://bit.ly/MSJobs\\_ProdQty\\_Supp](http://bit.ly/MSJobs_ProdQty_Supp)) ([http://bit.ly/MSJobs\\_IT\\_ProgMgr](http://bit.ly/MSJobs_IT_ProgMgr))

### SAN FRANCISCO, CA

**Program Managers:** (All levels, including Leads and Managers) Coordinate program development of computer software applications, systems or services working with development and product planning teams. ([http://bit.ly/MSJobs\\_ProgMgr](http://bit.ly/MSJobs_ProgMgr)) ([http://bit.ly/MSJobs\\_HW\\_ProgMgr](http://bit.ly/MSJobs_HW_ProgMgr)) ([http://bit.ly/MSJobs\\_ProdQty\\_Supp](http://bit.ly/MSJobs_ProdQty_Supp)) ([http://bit.ly/MSJobs\\_IT\\_ProgMgr](http://bit.ly/MSJobs_IT_ProgMgr))

### CAMBRIDGE, MA

**Program Managers:** (All levels, including Leads and Managers) Coordinate program development of computer software appli-

cations, systems or services working with development and product planning teams. ([http://bit.ly/MSJobs\\_ProgMgr](http://bit.ly/MSJobs_ProgMgr)) ([http://bit.ly/MSJobs\\_HW\\_ProgMgr](http://bit.ly/MSJobs_HW_ProgMgr)) ([http://bit.ly/MSJobs\\_ProdQty\\_Supp](http://bit.ly/MSJobs_ProdQty_Supp)) ([http://bit.ly/MSJobs\\_IT\\_ProgMgr](http://bit.ly/MSJobs_IT_ProgMgr))

### BOSTON, MA

**Program Managers:** (All levels, including Leads and Managers) Coordinate program development of computer software applications, systems or services working with development and product planning teams. ([http://bit.ly/MSJobs\\_ProgMgr](http://bit.ly/MSJobs_ProgMgr)) ([http://bit.ly/MSJobs\\_HW\\_ProgMgr](http://bit.ly/MSJobs_HW_ProgMgr)) ([http://bit.ly/MSJobs\\_ProdQty\\_Supp](http://bit.ly/MSJobs_ProdQty_Supp)) ([http://bit.ly/MSJobs\\_IT\\_ProgMgr](http://bit.ly/MSJobs_IT_ProgMgr))

### IRVING, TX

**Program Managers:** (All levels, including Leads and Managers) Coordinate program development of computer software applications, systems or services working with development and product planning teams. ([http://bit.ly/MSJobs\\_ProgMgr](http://bit.ly/MSJobs_ProgMgr)) ([http://bit.ly/MSJobs\\_HW\\_ProgMgr](http://bit.ly/MSJobs_HW_ProgMgr)) ([http://bit.ly/MSJobs\\_ProdQty\\_Supp](http://bit.ly/MSJobs_ProdQty_Supp)) ([http://bit.ly/MSJobs\\_IT\\_ProgMgr](http://bit.ly/MSJobs_IT_ProgMgr))

Multiple job openings are available for each of these categories. To view detailed job descriptions and minimum requirements, and to apply, visit the website address listed. EOE.

## TECHNICAL Oracle America, Inc.

has openings for

# TECHNICAL ANALYST

positions in

Colorado Springs, CO.

Job duties include: Deliver solutions to the Oracle customer base while serving as an advocate for customer needs. Offer strategic technical support to assure the highest level of customer satisfaction.

Apply by e-mailing resume to [valerie.mayes@oracle.com](mailto:valerie.mayes@oracle.com), referencing 385.19923.

Oracle supports workforce diversity.

**SENIOR BIG DATA ENGINEER**, General Motors, Detroit, MI. Use Hadoop platform &Kafka, HBASE, Hive &Spark to perform Big Data analysis inclgd complete vehicle prognostics &diagnostics. Design, &optimize infotainment Customer Relationship Mgmt, Provisioning &Vehicle Diagnostics Systems, &Billing &Revenue mgmt business processes. Mentor, train &lead internal customers &colleagues in dvlpg &deploying emerging technologies for business use cases. Perform Data Mgmt using InfoSphere Information Server Product Suite (inclgd Information Governance Catalog) for automotive telematics platforms. Perform Data Replication using Oracle Golden Gate &InfoSphere Change Data Delivery. Use large scale database platforms such as Hadoop, Teradata,

IBM Big Insights, &Oracle ExaData. Dvlp BI reports using Cognos for automotive telematics platforms. Install, configure Information Server suite on Redhat, Suse &Oracle Enterprise Linux servers. Master, Computer Science, Computer Science &Engrg, or Information Systems. 36 mos exp as Sr Application Architect, Solution Architect, Big Data Engr, or related, performing Data Mgmt using InfoSphere Information Server Product Suite for automotive telematics platforms, &Data Replication using Oracle Golden Gate &InfoSphere Change Data Delivery. Will accept bachelor's or foreign equiv degree, in Computer Science, Computer Science and Engineering, or Information Systems, followed by at least 5 yrs of progressive exp in the specialty, in lieu

of required education &exp. Will also accept any equally suitable combination of education, training, &/or experience which would qualify applicant to perform job offered. Mail resume to Ref#2251, GM Global Mobility, 300 Renaissance Center, MC:482-C32-D44, Detroit, MI 48265.

**ENGINEERING. PROJECT MANAGER 1.** CH2M HILL, Inc. in Kansas City, MO seeks a Project Manager 1 to lead and manage projects/programs that are aligned to the water business group markets and underlying business objectives. To apply, mail resume to: Shelly Saitta, CH2M HILL, 9191 S. Jamaica St., Englewood, CO 80112. Must reference job code: 13176.



**Microsoft**

Microsoft Corporation currently has the following openings (job opportunities available at all levels, e.g., Principal, Senior and Lead levels):

**ALISO VIEJO, CA**

**Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers):** Responsible for developing or testing computer software applications, systems or services. ([http://bit.ly/MSJobs\\_SDE](http://bit.ly/MSJobs_SDE)) ([http://bit.ly/MSJobs\\_IT\\_SDE](http://bit.ly/MSJobs_IT_SDE))

**LOS ANGELES, CA**

**Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers):** Responsible for developing or testing computer software applications, systems or services. ([http://bit.ly/MSJobs\\_SDE](http://bit.ly/MSJobs_SDE)) ([http://bit.ly/MSJobs\\_IT\\_SDE](http://bit.ly/MSJobs_IT_SDE))

**MOUNTAIN VIEW, PALO ALTO, SUNNYVALE, CA**

**Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers):** Responsible for developing or testing computer software applications, systems or services. ([http://bit.ly/MSJobs\\_SDE](http://bit.ly/MSJobs_SDE)) ([http://bit.ly/MSJobs\\_IT\\_SDE](http://bit.ly/MSJobs_IT_SDE))

**SAN FRANCISCO, CA**

**Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers):** Responsible for developing or testing computer software applications, systems or services. ([http://bit.ly/MSJobs\\_SDE](http://bit.ly/MSJobs_SDE)) ([http://bit.ly/MSJobs\\_IT\\_SDE](http://bit.ly/MSJobs_IT_SDE))

**BOSTON, MA**

**Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers):** Responsible for developing or testing computer software applications, systems or services. ([http://bit.ly/MSJobs\\_SDE](http://bit.ly/MSJobs_SDE)) ([http://bit.ly/MSJobs\\_IT\\_SDE](http://bit.ly/MSJobs_IT_SDE))

**Program Managers:** (All levels, including Leads and Managers) Coordinate program development of computer software applications, systems or services working with development and product planning teams. ([http://bit.ly/MSJobs\\_ProgMgr](http://bit.ly/MSJobs_ProgMgr)) ([http://bit.ly/MSJobs\\_HW\\_ProgMgr](http://bit.ly/MSJobs_HW_ProgMgr)) ([http://bit.ly/MSJobs\\_ProdQty\\_Supp](http://bit.ly/MSJobs_ProdQty_Supp)) ([http://bit.ly/MSJobs\\_IT\\_ProgMgr](http://bit.ly/MSJobs_IT_ProgMgr))

**CAMBRIDGE, MA**

**Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers):** Responsible for developing or testing computer software applications, systems or services. ([http://bit.ly/MSJobs\\_SDE](http://bit.ly/MSJobs_SDE)) ([http://bit.ly/MSJobs\\_IT\\_SDE](http://bit.ly/MSJobs_IT_SDE))

**MORRISVILLE, NC**

**Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers):** Responsible for developing or testing computer software applications, systems or services. ([http://bit.ly/MSJobs\\_SDE](http://bit.ly/MSJobs_SDE)) ([http://bit.ly/MSJobs\\_IT\\_SDE](http://bit.ly/MSJobs_IT_SDE))

**FARGO, ND**

**Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers):** Responsible for developing or testing computer software applications, systems or services. ([http://bit.ly/MSJobs\\_SDE](http://bit.ly/MSJobs_SDE)) ([http://bit.ly/MSJobs\\_IT\\_SDE](http://bit.ly/MSJobs_IT_SDE))

**NEW YORK, NY**

**Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers):** Responsible for developing or testing computer software applications, systems or services. ([http://bit.ly/MSJobs\\_SDE](http://bit.ly/MSJobs_SDE)) ([http://bit.ly/MSJobs\\_IT\\_SDE](http://bit.ly/MSJobs_IT_SDE))

**RESTON, VA**

**Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers):** Responsible for developing or testing computer software applications, systems or services. ([http://bit.ly/MSJobs\\_SDE](http://bit.ly/MSJobs_SDE)) ([http://bit.ly/MSJobs\\_IT\\_SDE](http://bit.ly/MSJobs_IT_SDE))

**REDMOND, WA**

**Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers):** Responsible for developing or testing computer software applications, systems or services. Requires domestic and international travel up to 25%. ([http://bit.ly/MSJobs\\_SDE](http://bit.ly/MSJobs_SDE)) ([http://bit.ly/MSJobs\\_IT\\_SDE](http://bit.ly/MSJobs_IT_SDE))

**Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers):** Responsible for developing or testing computer software applications, systems or services. Telecommuting permitted. ([http://bit.ly/MSJobs\\_SDE](http://bit.ly/MSJobs_SDE)) ([http://bit.ly/MSJobs\\_IT\\_SDE](http://bit.ly/MSJobs_IT_SDE))

**Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers):** Responsible for developing or testing computer software applications, systems or services. ([http://bit.ly/MSJobs\\_SDE](http://bit.ly/MSJobs_SDE)) ([http://bit.ly/MSJobs\\_IT\\_SDE](http://bit.ly/MSJobs_IT_SDE))

**Multiple job openings are available for each of these categories. To view detailed job descriptions and minimum requirements, and to apply, visit the website address listed. EOE.**

## CAREER OPPORTUNITIES

**LIGHTBEND SEEKS:** Sr. Software Engineer Full Stack -Architect, design, develop, deliver company products, platforms and API's, create automated tests/infrastructures; Present at conferences, meetups worldwide. 10% Domestic Travel. Worksite: May telecommute from any US location or headquarters (San Francisco, CA 94105). To apply send resume/cover(20170301HE) to Lightbend, 625 Market St., Ste. 1000, San Francisco, CA, 94105.

**ERICSSON INC.** has openings for the following positions: **ENGINEER-SERVICES SOFTWARE \_ OVERLAND PARK, KS** integration, certification, & implementation of 4G LTE technology into customer's lab facilities & first market appls. Travel req'd. Mail resume to Ericsson Inc. 6300 Legacy Dr, R1-C12, Plano, TX 75024 & indicate applying for **17-KS-2569. DATA SCIENTIST/DATA ANALYTICS ARCHITECT \_ PLANO, TX** provide analysis of big data, mathematical modeling, & Data Mining techniques applied to data monetization opportunities for telecomm. Travel req'd. Telecommuting available. Mail resume to

Ericsson Inc. 6300 Legacy Drive, R1-C12 Plano, TX 75024 and indicate applying for **16-TX-3320.**

**SOLUTIONS ARCHITECT ASCENSION HEALTH-IS, INC.** is seeking a Solutions Architect in Austin, Texas to identify issues between the architecture and vendor solution implementation; identify solutions and formulate implementation and migration strategies; and provide expertise in technical knowledge, methodology and framework, data integration and data flows, Enterprise Architecture, and software development and coding. Contact Denise D. Dodge, Project Lead, 11775 Borman Drive, St. Louis, MO 63146.

**VALASSIS COMMUNICATIONS INC.** has an opening for the position of Applications Engineer in WINDSOR, CT to analyze, write, test, & deploy software applications & research, design, document, & modify software specifications. To apply mail resume to Valassis, Attn: Patty [CT-11307.26], 19975 Victor Parkway, Livonia, MI 48152.

### TECHNICAL

## Oracle America, Inc.

has openings for

# TECHNICAL ANALYST

positions in **Lehi, UT.**

Job duties include: Deliver solutions to the Oracle customer base while serving as an advocate for customer needs.

Apply by e-mailing resume to  
Al.Rocek@oracle.com,  
referencing 385.19075.

Oracle supports workforce diversity.

### SOFTWARE

## Oracle America, Inc.

has openings for

# SOFTWARE DEVELOPER

positions in **Burlington, Massachusetts.**

Job duties include: design, develop, troubleshoot and/or test/QA software; as a member of the software engineering division, apply knowledge of software architecture to perform tasks associated with developing, debugging, or designing software applications or operating systems according to provided design specification.

Apply by e-mailing resume to  
rob.levine@oracle.com,  
referencing 385.18388.

Oracle supports workforce diversity.

### TECHNICAL

## Oracle America, Inc.

has openings for

# TECHNICAL ANALYST

positions in **Lehi, Utah.**

Job duties include: Deliver solutions to the Oracle customer base while serving as an advocate for customer needs.

Apply by e-mailing resume to  
don.debeaux@oracle.com,  
referencing 385.19975.

Oracle supports workforce diversity.

### SOFTWARE

## Oracle America, Inc.

has openings for

# SOFTWARE DEVELOPER

positions in **Minneapolis, Minnesota.**

Job duties include: design, develop, troubleshoot and/or test/QA software; as a member of the software engineering division, apply knowledge of software architecture to perform tasks associated with developing, debugging, or designing software applications or operating systems according to provided design specifications.

Apply by e-mailing resume to  
shanti.atrangada@oracle.com,  
referencing 385.18110.

Oracle supports workforce diversity.

**SOFTWARE**  
**Oracle America, Inc.**

has openings for

**SOFTWARE  
DEVELOPER**

positions in **Irvine, CA.**

Job duties include: Design, develop, troubleshoot and/or test/QA software. As a member of the software engineering division, apply knowledge of software architecture to perform tasks associated with developing, debugging, or designing software applications or operating systems according to provided design specifications. Travel to various unanticipated sites throughout the United States required.

Apply by e-mailing resume to  
john.melon@oracle.com,  
referencing 385.19985.

Oracle supports workforce diversity.

**SOFTWARE**  
**Oracle America, Inc.**

has openings for

**SOFTWARE  
DEVELOPER**

positions in **Seattle, Washington.**

Job duties include: design, develop, troubleshoot and/or test/QA software; as a member of the software engineering division, apply knowledge of software architecture to perform tasks associated with developing, debugging, or designing software applications or operating systems according to provided design specifications.

Apply by e-mailing resume to  
rob.showalter@oracle.com,  
referencing 385.20615.

Oracle supports workforce diversity.

**TECHNICAL**  
**Oracle America, Inc.**

has openings for

**TECHNICAL  
ANALYSTS-  
SUPPORT**

positions in **Colorado Springs, Colorado.**

Job duties include: offer strategic technical support to assure the highest level of customer satisfaction; quickly isolate and clarify customer issues, apply technical knowledge to reproduce product issues, identify solutions based on testing results and/or comprehensive research, all while meeting multiple customers' critical timelines.

Apply by e-mailing resume to  
valerie.mayes@oracle.com,  
referencing 385.19769.

Oracle supports workforce diversity.

**QA ANALYST**  
**Oracle America, Inc.**

has openings for

**QA  
ANALYST**

positions in **Bedford, Massachusetts.**

Job duties include: responsible for developing, applying and maintaining quality standards for company products with adherence to both internal and external standards; develop and execute software test plans. Analyze and write test standards and procedures.

Apply by e-mailing resume to  
vida.haririan@oracle.com,  
referencing 385.20197.

Oracle supports workforce diversity.

**SOFTWARE**  
**Oracle America, Inc.**

has openings for

**SOFTWARE  
DEVELOPER**

positions in **Orlando, FL.**

Job duties include: Design, develop, troubleshoot and/or test/QA software.

Apply by e-mailing resume to  
maher.muhanha@oracle.com,  
referencing 385.18437.

Oracle supports workforce diversity.

**TECHNICAL**  
**Oracle America, Inc.**

has openings for

**TECHNICAL  
ANALYST-  
SUPPORT**

positions in **Orlando, FL.**

Job duties include: deliver solutions to the Oracle customer base while serving as an advocate for customer needs; offer strategic technical support to assure the highest level of customer satisfaction.

Apply by e-mailing resume to  
dana.downing@oracle.com,  
referencing 385.20037.

Oracle supports workforce diversity.



**Juniper Networks is recruiting for our Sunnyvale, CA office:**

**Engineering Program Manager #50766:** Design, develop and monitor the execution of software development projects in partnership with engineering team.

**Software Engineer Staff #37399:** Analyze complex business problems and collaborate with project team members to ensure development and implementation of new data driven business solutions.

**Test Engineer #34383:** Design, develop script and implement sophisticated tests to verify software feature functionality for Company products.

**Software Engineer #23245:** Design, develop troubleshoot and debug detailed software functional and design specifications for new software features on Juniper products.

**Software Engineer Senior Staff #4587:** Design, develop, troubleshoot and debug high performance forwarding software for current generation and next generation Terabit Core Routers.

**Systems Engineer #27726:** Assist Company Account Managers with the development and delivery of technical solutions and sales activities for new and existing opportunities. Telecommuting allowed.

**Software Engineer #38704:** Design, develop, troubleshoot and debug component and system level diagnostics software/firmware using Python programming language.

**Software Engineer Staff #6201:** Design, develop, troubleshoot and debug design specifications for network security products in physical and virtual environments that are being implemented in Data Center, Mobile LTE, Enterprise and Cloud deployments.

**Software Engineer #37060:** Analyze, design, develop, program, debug and modify test management and workflow related software applications for JUNOS developers and test communities.

**QA Engineer #39696:** Work within a security system testing team to debug and test complex multi-service networking and network security products for Data Center and cloud.

**Software Engineer Staff #28235:** Participate in new product development and contribute technical innovation to Juniper network security platform. Design and develop solutions for new features and enhancements for existing Juniper security products.

**Software Engineer #35871:** Design, develop, troubleshoot and debug high-performance packet forwarding software for the current generation and next generation terabit-core routers.

**Hardware Engineer #30838:** Plan, design and develop signal integrity design rules for boards, packages and ASICs with high speed electrical interfaces.

**Technical Systems Analyst #38746:** Analyze, design, develop and support SAP technical enhancements and deliver new and complex business solutions by integrating modules of SAP and external systems to meet business and user requirements.

**Sales Demonstration Engineer #26725:** Work with customers, Company business units and account teams to demonstrate Company products and technologies through Proof of Concept (POC) testing.

**Software Engineer #29261:** Design, develop, troubleshoot and debug functional and design specifications for software products and enhancement.

**Systems Engineer #29195:** Design, develop and test software modules to support security and networking features in Company products. Telecommuting allowed. Travel required.

**Software Engineer Sr. Staff #8050:** Work with product line management and customers to analyze, design, and modify device and network manageability requirements and help implement solutions.

**Test Engineer Staff #15855:** Using product definitions, scaling targets, and customer use scenarios, design, develop and implement product performance and scaling tests for high performance Ethernet switching platforms.

**Test Engineer #26482:** Provide system-level testing for routing protocols

and SDN. Develop and execute test plans for new functionality and features developed in Junos.

**Functional Systems Analyst Staff #38661:** Design and document solutions to address complex business problems. Analyze business processes, gather project requirements, and contribute to design and review workshops.

**Software Engineer Staff #14818:** Design, develop, troubleshoot and debug software modules for Company's products and platforms.

**Software Engineer #37443:** Work closely with product marketing software and hardware development teams to review design specification, then develop functional test plan for review.

**Systems Engineer #11815:** Develop and deliver technical sales strategy and solutions by providing pre-sales technical support, training, and lab support to meet end customer requirements. Telecommuting allowed.

**Hardware Engineer #30222:** Responsible for the testing and characterization of optical sub-systems and modules used in Company platforms to evaluate reliability of materials, properties and techniques used in production.

**Software Engineer #39619:** Design, develop, troubleshoot and debug ingestion pipelines for a variety of data from enterprise systems, routers, switches, and firewalls.

**Software Engineer #18288:** Analyze, design, program, debug and modify novel operating system techniques and algorithms for On-Device management software products.

**Sr. Strategic Delivery Engineer #28167:** Design, develop and implement testing efforts which could include automated testing, regression testing, scale testing, migration testing, certification testing and end to end solution testing for Network security products.

**Software Engineer #29835:** Develop and maintain virtual machine creation and control software. Debug issues found within the virtual cloud environment. Implement new OpenStack and Contrail based functionality.

**Software Engineer #39316:** Design and

develop applications and web services for JUNOS developers and test teams. Research software technologies to optimize existing infrastructure or software design.

**Software Engineer #7391:** Design, develop, troubleshoot and debug functional, design and architectural specifications for next generation interfaces and long haul optical products.

**Test Engineer #39241:** Design, develop and implement testing for complex multi-service networking and network security products.

**Product Marketing Manager #29873:** Create and maintain marketing assets for Company product lines and solutions across all media types, to establish, enhance, and distinguish product within industry.

**Sustaining Engineer #51117:** Diagnose, troubleshoot, repair and debug complex networking issues in the Data Center and Switching areas.

**Software Engineer #36797:** Design, develop, troubleshoot, debug and modify large scale networking software.

**Juniper Networks is recruiting for our Westford, MA office:**

**Technical Support Engineer #37316:** Provide technical phone support to customers and assistance to other support engineers to resolve product issues.

**Juniper Networks is recruiting for our Herndon, VA office:**

**Systems Engineer #33983:** Provide onsite technical guidance to the customers on Juniper's routing, switching, security solutions and services. Understand customers' environment and position Juniper's solutions within the customers' infrastructure. Must be willing to travel to client sites locally.

**Mail single-sided resume with job code # to  
Juniper Networks  
Attn: MS A.4.435  
1133 Innovation Way  
Sunnyvale, CA 94089**

**COMPUTER PROGRAMMER, MEMPHIS, TENNESSEE:** Limited domestic travel and/or occasional relocation to client sites nationwide to write code, forms, and script using Java, JSP, Servlets, JDBC, PL/SQL, HTML. Debug, troubleshoot existing code. Reply to: SVS Technologies Limited, 8700 Trail Lake Drive, #228, Memphis, TN 38125

**REPORTS DEVELOPER:** design & develop data models using Cognos; analyze & troubleshoot Cognos reporting. BS+5 yrs. or MS+3. marcela.niemeyer@laureate.net w/ Job # 20501BR in subj line. Job in McLean, Virginia. Laureate Education, Inc. EOE.

**MANAGER.** Job location: Miami, FL & any other unanticipated locations in U.S. Travel Required. Duties: Resp. for advising & implement. strategic solutions for Finance & Enterprise Performance Mgmt. (EPM) processes incl. Financial Close & Reporting, Acct. Reconciliation, Planning, Budgeting, Forecasting & Data Integration using Oracle EPM Products. Resp. for developing end-to-end

solution & maint. of Hyperion Financial Mgmt. (HFM) appls. Execute coding, testing & maint. activities. Provide critical appl. design recommend. to client which aligns with its long & short term IT & business strategies & build solutions as per external SEC reporting & mgmt. reporting requirements. Resp. for developing solutions for Hyperion FDMEE, Hyperion Acct. Reconciliation & Hyperion Planning appls. Advise clients on financial best practices & drive day-to-day execution of client engagements related to EPM project. Create strategy, blueprint, & process & data optimization for clients for Oracle EPM products. Requires: M.S. degree in Finance, Bus. Adm., or related field & 3 yrs. exp. in job offered or 3 yrs. exp. as a Consultant. Will accept B.S. (or foreign equiv.) & 5 yrs. exp. in the finance in lieu of M.S. & 3 yrs. exp. Concurrent exp. must incl.: 3 yrs. exp. executing coding, testing & maint. activities. Send resume (no calls) to: Michelle Ramirez, The Hackett Group, Inc., 1001 Brickell Bay Dr., Suite 3000, Miami, FL 33131.

**OCEAN PREDICTION  
POSTDOCTORAL POSITIONS  
Naval Research Laboratory,  
Stennis Space Center, MS**



The Naval Research Laboratory is seeking postdoctoral researchers to push forward the frontiers of ocean forecasting. cover a wide scope of physics including surface waves, thermohaline circulation, nearshore circulation, and ocean/atmosphere coupling from global to nearshore scales. This challenging work includes processing and analysis of satellite and in water observations, construction of numerical model systems on high performance computing and assimilation for predicting the ocean environment. For a quick overview of some of the research projects within the NRL oceanography division at Stennis Space Center, visit the web site: <http://www7320.nrlssc.navy.mil/projects.php>

Applicants must be a US citizen or permanent resident at time of application. Applications will be accepted until positions are filled. Please e-mail a resume and description of research interests:

**Gregg Jacobs: [jacobs@nrlssc.navy.mil](mailto:jacobs@nrlssc.navy.mil)**

## CAREER OPPORTUNITIES

**INFOSYS LIMITED** is in need of individuals to work full-time in Plano, Texas and various and unanticipated locations throughout the U.S. Must be willing to work anywhere in the U.S. as all job opportunities may involve relocation to various and unanticipated client site locations; any travel and/or relocation to be paid by employer pursuant to internal policy. We have multiple openings for each job opportunity, and are an Equal Opportunity Employer M/F/D/V. Please apply to Infosys Limited online at: <https://www.infosys.com/careers/job-opportunities/pages/index.aspx>. Select the box labeled "Technical" which will bring you to the site with the "Experienced Professionals" link in the right column. Click on the "Experienced Professionals" link. Once a user account has been created, please follow the link for 'Search Openings' and click on 'Advanced Search.' Enter the Req ID listed below in the "Auto Req ID" box. **Engagement Manager(s)** needed to navigate IT business account to identify different kinds of deals; form and lead pursuit teams; help pursuit

teams with customer context, competitor and industry context, details on pain-points and client introductions required for opening diverse service offerings in own account(s). Travel required. **(Req ID: 20548BR)**. **Industry Principal(s)** - U.S. needed to

lead IT change management tracks for medium to large engagements. Will lead the engagement effort at different stages of IT projects from business process consulting, and problem definition to solution design and deployment. Travel required. **(Req ID: 20549BR)**.



### TECHNOLOGY Oracle America, Inc.

has openings for the following positions (all levels/types) in **San Mateo County**, including **Redwood Shores, CA** and **San Bruno, CA**; **Alameda County**, including **Pleasanton, CA**; **San Francisco, CA**; **Santa Clara County**, including **Santa Clara** and **San Jose, CA**; and other locations in the **San Francisco Bay Area**. Some positions may allow for telecommuting.

**Hardware Developers (HWD317):** Evaluate reliability of materials, properties and techniques used in production; plan, design and develop electronic parts, components, integrated circuitry, mechanical systems, equipment and packaging, optical systems and/or DSP systems.

**Product Managers (PM317):** Participate in all software and/or hardware product development life cycle activities. Move software products through the software product development cycle from design and development to implementation, testing, and/or marketing.

**Software Developers (SWD317):** Design, develop, troubleshoot and/or test/QA software.

**Applications Developers (APD317):** Analyze, design, develop, troubleshoot and debug software programs for commercial or end user applications. Write code, complete programming and perform testing and debugging of applications.

**Programmer Analysts (PA317):** Analyze user requirements to develop, implement, and/or support Oracle's global infrastructure.

**Technical Analysts-Support (TAS317):** Deliver solutions to the Oracle customer base while serving as an advocate for customer needs. Offer strategic technical support to assure the highest level of customer satisfaction.

**Consultants (TCONS317):** Analyze requirements and deliver functional and technical solutions. Implement products and technologies to meet post-sale customer needs. Travel to various unanticipated sites throughout the U.S. required.

**Sales Consultants (TSC317):** Provide presales technical/functional support to prospective customers. Design, validate and present Oracle's software solutions to include product concepts and future direction. Travel to various unanticipated sites throughout the U.S. required.

**Software Developers (TSWD317):** Design, develop, troubleshoot and/or test/QA software. Travel to various unanticipated sites throughout the U.S. required.

**Applications Developers (TAPD317):** Analyze, design, develop, troubleshoot and debug software programs for commercial or end user applications. Write code, complete programming and perform testing and debugging of applications. Travel to various unanticipated sites throughout the U.S. required.

Submit resume to [applicant\\_us@oracle.com](mailto:applicant_us@oracle.com). Must include job#. Oracle supports workforce diversity.

**CLASSIFIED LINE AD SUBMISSION DETAILS:** Rates are \$425.00 per column inch (\$640 minimum). Eight lines per column inch and average five typeset words per line. Send copy at least one month prior to publication date to: Debbie Sims, Classified Advertising, *Computing Edge* Magazine, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720; (714) 816-2138; fax (714) 821-4010. Email: dsims@computer.org.

In order to conform to the Age Discrimination in Employment Act and to discourage age discrimination, *Computing Edge* may reject any advertisement containing any of these phrases or similar ones: "...recent college grads...", "...1-4 years maximum experience...", "...up to 5 years experience," or "...10 years maximum experience." *Computing Edge* reserves the right to append to any advertisement without specific notice to the advertiser. Experience ranges are suggested minimum requirements, not maximums. *Computing Edge* assumes that since advertisers have been notified of this policy in advance, they agree that any experience requirements, whether stated as ranges or otherwise, will be construed by the reader as minimum requirements only. *Computing Edge* encourages employers to offer salaries that are competitive, but occasionally a salary may be offered that is significantly below currently acceptable levels. In such cases the reader may wish to inquire of the employer whether extenuating circumstances apply.



香港中文大學  
The Chinese University of Hong Kong

Applications are invited for:-

**Faculty of Engineering  
Professors / Associate Professors / Assistant Professors  
(Ref. 1700004N)**

The Faculty of Engineering is seeking several faculty posts at Professor / Associate Professor / Assistant Professor levels with prospect for substantiation. The professors will play a significant role in the Cyber Security Center, which will be established by the Faculty of Engineering.

Cyber security is identified as one of the Faculty's strategic research areas, to be developed by both the Department of Computer Science & Engineering and Department of Information Engineering. Talented candidates are sought to complement existing efforts and create new synergies. Candidates in the following areas are encouraged to apply:

- cryptography and computational theory in security
- network, system and software security
- data security and privacy
- computer forensic
- hardware and IoT security

Applicants should have a relevant PhD degree and a good scholarly record demonstrating potential for teaching and research excellence.

Appointments will normally be made on contract basis for up to three years initially commencing August 2017, which, subject to performance and mutual agreement, may lead to longer-term appointment or substantiation later. The exact start date can be worked out with the successful applicants.

Applications will be accepted until the posts are filled.

**Application Procedure**

Applicants please upload the full resume with a cover letter, copies of academic credentials, publication list with abstracts of selected published papers, a research plan, a teaching statement, together with names and e-mails addresses of three to five referees to whom the applicant's consent has been given for their providing reference (unless otherwise specified).

The University only accepts and considers applications submitted online for the posts above. For more information and to apply online, please visit <http://career.cuhk.edu.hk>.

**TECHNOLOGY**

# Intuit Inc.

has openings for the following positions in **Mountain View, California:**

**Staff Application Operations Engineers (Job code: I-292):** Apply a full understanding of the business, the customer, and the solutions that a business offers to effectively design, develop, and implement operational capabilities, tools and processes. **Senior Business Data Analysts (Job code: I-2311):** Write SQL queries to pull sales data using Site Catalyst.

**Technical Data Analysts (Job code: I-2422):** Partner with marketing, finance, analytics, and cross-functional teams to interpret large volumes of data, address key business questions, from hypothesis to execution.

Positions in **Woodland Hills, California:**

**Software Engineers (Job code: I-3134):** Apply software development practices to design, implement, and support individual software projects.

To apply, submit resume to Intuit Inc., Attn: Olivia Sawyer, J203-6, 2800 E. Commerce Center Place, Tucson, AZ 85706.

You must include the job code on your resume/cover letter. Intuit supports workforce diversity.

## Apple Inc. has the following job opportunities in Cupertino, CA:

**Software Engineer Applications (Req# 9KSNDX)** Dev the next gen geolocation system. Optimize & improve processing pipeline.

**Software Development Engineer (Req# 9U4RR8)** (Lang Req – German) Cond sw qual test to ensure the qual of groundbreaking tech for large scale sys, spoken lang, big data, and AI with a focus on the German user exp.

**Software Engineer Applications (Req# A6F33Z)** Des, build, and sup new, critical infra systems and frameworks.

**Software QA Engineering Lead (Req# 9X5SQN)** Lead Web QA Team. Assist & guide in dev javascript-based web automation.

**Software Engineer Applications (Req# 9WNPT9)** Des & dev SW for telemetry systems.

**Software Engineer Applications (Req# A9GTLG)** Des & drive new prods/pltfms frm concept to implem in the Architecture tm.

**Software Engineer Applications (Req# 9XE23J)** Res, des, dev, implmnt, & debug iOS apps. Travel req'd 10%.

**Software Engineer Applications (Req# 9F8TLF)** Research, des, dev, & productize machine learning mdl for iTunes & App Store recommender sys for personalization.

**Software Development Engineer (Req# A2GTX4)** Dsgn & implmnt features in iAd Workbench News Web App & API.

**Software Engineer Applications (Req# 9XJVVXU)** Des and dev large scale distributed sys to sup workflow engines that process and compute large amt of data in var domains.

**Software Engineer Applications (Req# 9WJ2EP)** Develop REST Web services for third party interactions.

**Software Quality Assurance Engineer (Req# A9CU96)** Build app test cases (auto & manual) & exec req'd tests.

**Software Development Engineer (Req# 9Q5MF8)** Des & dev SW testing for camera, video, audio, & graphics of Apple prods.

**Software Engineer Applications (Req# ACQ24M)** Dsgn, dev, & support Apple's realtime cmmunications infrstrctre.

**Software Development Engineer (Req# AAZPJ3)** Perf diff types of test ranging from fnctnl, feature, regression & auto test.

**Software Engineer Applications (Req# AHYU5W)** Validate sys capacity to ensure suff resource availability.

**Software Development Engineer (Req# A8XSM7)** Dsgn & dvlp real-time comm SW for FaceTime.

**Software Development Manager (Req# 9DTNVH)** Des & lead dev of Siri SW & languags.

**Software Engineer Applications (Req# 9PNSHN)** Des & implemnt SW to aid in the anlysis & correctn of client-facng prdcts, usng machine learnng technques, on a globl scale.

**Software Development Engineer (Req# A3FW45)** Des, dev, debug & port bootloader & driver SW for mobile realtime OS.

**Software Development Engineer (Req# A7A24T)** Dsgn, dvlp & implmnt search relevance & quality sys

**Software Development Engineer (Req# A3RW5F)** Dvlp the SW build sys. Set up & supprt CI & dplymnt tools.

**Software Development Engineer (Req# 9QEP8Z)** (Lang Req – Span) Conduct SW Quality Testing to ensure the quality of groundbreaking tech for lrg scale sys, spoken lang,

big data, & artificial intelligence w/ focus on Span user exp.

**Software Engineer Applications (Req# A7T2TW)** Des & dev SW for iCloud reporting sys.

**Systems Design Engineer (Req# 9CYVYJ2ND)** Des and dev int SW tools and sys for the Wireless Des dept.

**Mechanical Quality Engineer (Req# 9T65WU)** Support new Apple product launches from a mech, tooling, eng & product dev level. Travel req 30%.

**Product Design Engineer (Req# 9TATXU)** Dev & implement high precision factory & manufacturing automation processes in high volume environment for Apple products. Travel required 35%.

**Product Quality Engineer (Req# 99N36Z)** Dev & impl new tech to enhance production quality, capacity & efficiency. Travel req 30%.

**Software Development Engineer (Req# 9UJVV5)** Des, dev, debug & test Apple Maps SW on iOS & OSX.

**Software Quality Assurance Engineer (Req# A7449N)** Drive SW qual & integration to improve cust exp.

**Software Engineer Systems (Req# 9TU28B)** Dev SW to automate engineering data collection in the RF Hardware organizatn.

**Software Development Manager (Req# 9E5QEN)** Mange the des of algorithms tht cntrl a cmplx mchncl/ elctrcal sys.

**Development Engineer (Req# A8R2LR)** Dsgn, implmnt & maintain highly scure & redundant ntwrk envirmnts.

**Software QA Engineer (Req# A5EVZJ)** Dsgn, dev, excute & analyze low lvl automated tstts fr mbile devcs.

**ASIC Design Engineer (Req# AF-H2UQ)** Archtct, dsgn & vrfy Pwr Mangmnt ICs.

**Software Engineer Applications (Req# A6D8PM)** Dsgn & dvlp mbile & web apps w automated tst.

**Software Engineer Systems (Req# 9T8VNY)** Archtct, spcfy, dsgn, dvlp, & launch Apple's touch tech prdct chrctriztion & prdctn instrmntation SW. Travel Req'd 25%

**Systems Design Engineer (Req# 9U32VA)** Dev & optimize RF automation sys used on Apple's newest prods incl iPhones, iPods, iPads & other wireless tech. Travel req'd: 20%.

**Software Quality Assurance Engineer (Req# 9TA3Q6)** Test audio & vid tech w manual & automation tst methods.

**Software Development Engineer (Req# 9Y6SJE)** Des & dev SW and test cases for testing map SW.

**Software Quality Assurance Engineer (Req# A663S5)** Des test plans to test wirless SW & HW featrs.

**Software Development Engineer (Req# A393BM)** Perform sys bring-up of new HW platforms.

**Hardware Development Engineer (Req# 9KT2B8)** Dsgn antennas for mobile communication devices. Travel req. 15%.

**Software Development Engineer (Req# 9WY3JX)** Qualify Wi-Fi feature fcnctnalty & prfrmnce acrss various Apl pltfms.

**Hardware Development Engineer (Req# A8SVW4)** Dev motion & behr predictn algs & corrspndng SW in C++/MATLAB.

**Development Engineer (Req# 9R2NK7)** Detct & rspnd to in-formtn secrt'y threats.

**Localization Engineer (Technical Translator) (Req# 9FM2U6)** Translate Apple SW prdcts & rel materials from English to Simplified Chinese.

**Software Development Engineer (Req# 9VSSMZ)** Rspnsble for the wireless quality on Apple macOS, iOS, WatchOS, & TVOS devices.

**Hardware Development Engineer (Req# A7M53A)** Analyze, debug, & imprv perf for new comp sys under dev.

**Engineering Project Lead (Req# A4YU2G)** Plan, coordnte, & rept the status of SW eng prjcts.

**Software Development Engineer (Req# 9FJTY5)** Des, impl & debug location frmwrk SW, APIs & functionality on iOS, macOS, watchOS & tvOS.

**Software Development Engineer (Req# A7245T)** Dsgn & implmnt features for contnt mngmnt app for lrge nmbr of pblshrs.

**Research Scientist/Engineer (Req# A6SN4E)** Anlyze data fr Apple's SW & srvcis. Write code to anlyz & visualze data.

**Software Development Engineer (Req# 9QJ38L)** Dev & maintain test plans. Carry out testing across a variety of netwrk topologies to ensure the hghst vid & audio quality.

**Firmware Engineer (Req# 9LUTLF)** Des & dev firmware for wireless audio prdcts.

**Software Development Engineer (Req# A9VQE8)** Analyze reqs, des & implement new user-facing web apps

**Software Development Engineer (Req# 9YNM8P)** Des & dvlp highly scalable distributed sys to process large scale data.

**Hardware Development Engineer (Req# 9LTQ3L)** Dsgn HW for Apple's Input Devices group. Travel Req 20%.

**Software Engineer Applications (Req# 9JG29L)** Des & dev SW for Apple News ecosys & feats pertaining to content mgmnt & authoring.

**Hardware Development Engineer (Req# 9SZ2Z2)** Dev, prototype & implmnt cutting-edge color imaging sol's & tech for Apple prods.

**Software Engineer Applications (Req# 9W9SUA)** Eval quality of algorithmic results in Apple Maps.

**Machine Learning Manager (Req# 9USPXF)** Dsgn & execute a Decision Sci prgrm fcsd on IDing, trgtng, & seizing fnc'l opportunities for AppleCare.

## Apple Inc. has the following job opportunities in Austin, TX:

**ASIC Design Engineer (Req# 9FT-JLV)** Deliver low-power analysis, optimization, estimation, & micro-archtcture power modeling for Apple's GPU.

**ASIC Design Engineer (Req# 9XJVYN)** Dsgn, simulate & verfy analog crcuits usng Cadence tools.

Refer to Req# & mail resume to Apple Inc.,  
ATTN: D.W.

1 Infinite Loop 104-1GM,  
Cupertino, CA 95014.

Apple is an EOE/AA m/f/  
disability/vets.

## TECHNOLOGY

# LinkedIn Corp.

LinkedIn Corp. has openings in our **Mtn View, CA** location for:

**Software Engineer (All Levels/Types) (SWE0317SV)** Design, develop & integrate cutting-edge software technologies; **Senior Software Engineer in Test (Test Engineer) (6597.1766)** Design & develop advanced test suites using object-oriented methodologies; **Manager, Software Engineering (6597.722)** Act as thought-leader for team and others, proposing and building investments in automation frameworks and tools to drive improvements and efficiencies across the larger test engineering organization; **Performance Engineer (6597.1655)** Design, develop & integrate cutting-edge software technologies; **User Experience Designer (6597.1193)** Develop original visual designs against specific communication priorities, utilizing tools such as Adobe Photoshop, Adobe Illustrator, Apple Keynote, Microsoft PowerPoint, Microsoft Word & online media collections like Getty Images & iStockPhoto; **Sr. Engineer, Site Reliability (6597.1050)** Serve as a primary point responsible for the overall health, performance & capacity of one or more Internet-facing service; **Senior Product Manager (6597.1896)** Analyze the competitive environment, customers & product metrics to determine the right set of features to drive engagement & usage on LinkedIn; **Engineer, Service Operations (6597.872)** Coordinate short- & long-term initiatives with LinkedIn DevOps teams, as well as partners (CDN, DNS, monitoring/measurement) – prioritizing & driving project closure on all sides; **Senior User Experience Designer (6597.1223)** Create holistic design solutions that address business, brand, and user requirements; **Associate Web Developer (6597.968)** Design, develop & integrate cutting-edge software technologies; **Product Manager, ProFinder (6597.1927)** Develop a comprehensive product & marketing roadmap to deliver business goals; **Manager, Software Engineering (6597.115)** Manage the performance & career development of a small team of engineers, & own significant parts of LinkedIn products that require design, architecture, & coding; **Quality Assurance Manager, Sales Systems (6597.1779)** Develop a QA program strategy which includes planning, evaluating, designing, architecting, implementing & maintaining test frameworks.

LinkedIn Corp. has openings in our **San Francisco, CA** location for:

**Software Engineer (All Levels/Types) (SWE0317SF)** Design, develop & integrate cutting-edge software technologies.

LinkedIn Corp. has openings in our New York, NY location for:

**Software Engineer (All Levels/Types) (SWE0317NY)** Design, develop & integrate cutting-edge software technologies; **Director, Software Engineering (6597.255)** Architect, design, develop, & support the most visible Internet-scale products & infrastructure at LinkedIn.

Please email resume to: [6597@linkedin.com](mailto:6597@linkedin.com). Must ref. job code above when applying.



# CONFERENCES

## *in the Palm of Your Hand*

IEEE Computer Society's Conference Publishing Services (CPS) is now offering conference program mobile apps! Let your attendees have their conference schedule, conference information, and paper listings in the palm of their hands.

The conference program mobile app works for **Android** devices, **iPhone**, **iPad**, and the **Kindle Fire**.

For more information please contact [cps@computer.org](mailto:cps@computer.org)





**TECHNICAL**  
**Oracle America, Inc.**  
 has openings for  
**TECHNICAL ANALYST-SUPPORT**  
 positions in **Lehi, Utah.**

Job duties include: deliver solutions to the Oracle customer base while serving as an advocate for customer needs; offer strategic technical support to assure the highest level of customer satisfaction. Travel to various unanticipated sites throughout the United States required.

Apply by e-mailing resume to [harshal.patil@oracle.com](mailto:harshal.patil@oracle.com), referencing 385.19904.  
 Oracle supports workforce diversity.

**TECHNICAL**  
**Oracle America, Inc.**  
 has openings for  
**TECHNICAL ANALYST**  
 positions in **Orlando, FL.**

Job duties include: Deliver solutions to the Oracle customer base while serving as an advocate for customer needs.

Apply by e-mailing resume to [Al.Rocek@oracle.com](mailto:Al.Rocek@oracle.com), referencing 385.19786.  
 Oracle supports workforce diversity.

**TECHNOLOGY**  
**Oracle America, Inc.**  
 has openings for  
**PRODUCT SUPPORT MANAGER**  
 positions in **Colorado Springs, CO.**

Job duties include: Contribute to the development of customer support management model and consistent business practices including team readiness to support new product releases or functionality.

Apply by e-mailing resume to [jon.green@oracle.com](mailto:jon.green@oracle.com), referencing 385.19495. Oracle supports workforce diversity.

**TECHNICAL**

**Oracle America, Inc.**

has openings for

**TECHNICAL  
ANALYSTS-  
SUPPORT**

positions in **Lehi, Utah.**

Job duties include: Deliver post-sales support and solutions to the Oracle customer base while serving as an advocate for customer needs.

Apply by e-mailing resume to  
vandana.sharma@oracle.com,  
referencing 385.19669.

Oracle supports workforce diversity.

**TECHNOLOGY**

**Oracle America, Inc.**

has openings for

**HARDWARE  
DEVELOPERS**

positions in **Burlington, MA.**

Job duties include: Evaluate reliability of materials, properties and techniques used in production; plan, design and develop electronic parts, components, integrated circuitry, mechanical systems, equipment and packaging, optical systems and/or DSP systems.

Apply by e-mailing resume to  
adam.reed@oracle.com,  
referencing 385.14655.

Oracle supports workforce diversity.



**Are Enemy Hackers Slipping  
through Your Team's Defenses?**

**Protect Your Organization from Hackers  
by Thinking Like Them**

**Take Our E-Learning Courses  
in the Art of Hacking**

You and your staff can take these courses where you are and at your own pace, getting hands-on, real-world training that you can put to work immediately.

[www.computer.org/artofhacking](http://www.computer.org/artofhacking)



Recognizing Excellence in High-Performance Computing

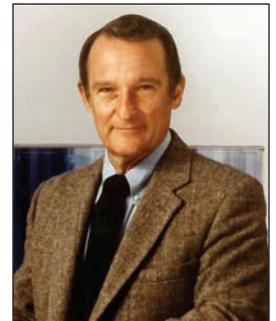
Nominations are Solicited for the

# SEYMOUR CRAY SIDNEY FERNBACH & KEN KENNEDY AWARDS

### SEYMOUR CRAY COMPUTER ENGINEERING AWARD

Established in late 1997 in memory of Seymour Cray, the Seymour Cray Award is awarded to recognize innovative contributions to high-performance computing systems that best exemplify the creative spirit demonstrated by Seymour Cray. The award consists of a crystal memento and honorarium of US\$10,000. This award requires 3 endorsements.

Sponsored by: IEEE  computer society



### SIDNEY FERNBACH MEMORIAL AWARD

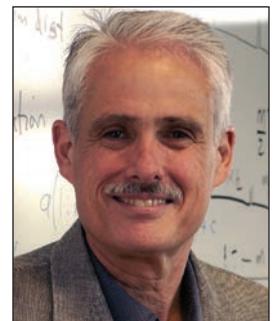
Established in 1992 by the Board of Governors of the IEEE Computer Society, this award honors the memory of the late Dr. Sidney Fernbach, one of the pioneers on the development and application of high-performance computers for the solution of large computational problems. The award, which consists of a certificate and a US\$2,000 honorarium, is presented annually to an individual for “an outstanding contribution in the application of high-performance computers using innovative approaches.” This award requires 3 endorsements.

Sponsored by: IEEE  computer society

### ACM/IEEE-CS KEN KENNEDY AWARD

This award was established in memory of Ken Kennedy, the founder of Rice University’s nationally ranked computer science program and one of the world’s foremost experts on high-performance computing. A certificate and US\$5,000 honorarium are awarded jointly by the ACM and the IEEE Computer Society for outstanding contributions to programmability or productivity in high-performance computing together with significant community service or mentoring contributions. This award requires 2 endorsements.

Cosponsored by: IEEE  computer society  Association for Computing Machinery



**Deadline: 1 July 2017**

All nomination details available at [awards.computer.org](http://awards.computer.org)



# NEW MEMBERSHIP OPTIONS FOR A BETTER FIT.



PREFERRED PLUS



TRAINING & DEVELOPMENT



RESEARCH



BASIC



STUDENT

## And a better match for your career goals.

IEEE Computer Society lets you choose your membership — and the benefits it provides — to fit your specific career needs. With four professional membership categories and one student package, you can select the precise industry resources, offered exclusively through the Computer Society, that will help you achieve your goals.

Learn more at [www.computer.org/membership](http://www.computer.org/membership).

IEEE  computer society

# Achieve your career goals with the fit that's right for you.

Explore your options below.



Select your membership	Preferred Plus		Training & Development		Research		Basic		Student
	\$60 IEEE Member	\$126 Affiliate Member	\$55 IEEE Member	\$115 Affiliate Member	\$55 IEEE Member	\$115 Affiliate Member	\$40 IEEE Member	\$99 Affiliate Member	\$8 Does not include IEEE membership
Computer magazine (12 digital issues)*	✓		✓		✓		✓		✓
ComputingEdge magazine (12 issues)	✓		✓		✓		✓		✓
Members-only discounts on conferences and events	✓		✓		✓		✓		✓
Members-only webinars	✓		✓		✓		✓		✓
Unlimited access to <i>Computing Now</i> , computer.org, and the new mobile-ready myCS	✓		✓		✓		✓		✓
Local chapter membership	✓		✓		✓		✓		✓
Skillsoft's Skillchoice™ Complete with 67,000+ books, videos, courses, practice exams and mentorship resources	✓		✓						✓
Books24x7 on-demand access to 15,000 technical and business resources	✓		✓						✓
Two complimentary Computer Society magazine subscriptions	✓				✓				
myComputer mobile app	30 tokens				30 tokens				30 tokens
Computer Society Digital Library	12 FREE downloads		Member pricing		12 FREE downloads		Member pricing		Included
Training webinars	3 FREE webinars		3 FREE webinars		Member pricing		Member pricing		Member pricing
Priority registration to Computer Society events	✓								
Right to vote and hold office	✓		✓		✓		✓		
One-time 20% Computer Society online store discount	✓								

\* Print publications are available for an additional fee. See catalog for details.

[www.computer.org/membership](http://www.computer.org/membership)

# IEEE Computer Society 2017 Call for **MAJOR AWARD NOMINATIONS**

*Help Recognize Computing's Most Prestigious*



IEEE Computer Society awards recognize outstanding achievements and highlight significant contributors in the teaching and R&D computing communities. All members of the profession are invited to nominate individuals who they consider most eligible to receive international recognition of an appropriate society award.

## **Computer Entrepreneur Award**

### **Sterling Silver Goblet**

Vision and leadership resulting in the growth of some segment of the computer industry.

## **Technical Achievement Award**

### **Certificate/\$2,000**

Contributions to computer science or computer technology.

## **Harry H. Goode Memorial Award**

### **Bronze Medal/\$2,000**

Information sciences, including seminal ideas, algorithms, computing directions, and concepts.

## **Hans Karlsson Award**

### **Plaque/\$2,000**

Team leadership and achievement through collaboration in computing standards.

## **Richard E. Merwin Award for Distinguished Service**

### **Bronze Medal/\$5,000**

Outstanding volunteer service to the profession at large, including service to the IEEE Computer Society.

## **Harlan D. Mills Award**

### **Plaque/\$3,000**

Contributions to the practice of software engineering through the application of sound theory.

## **Computer Pioneer Award**

### **Silver Medal**

Pioneering concepts and development of the computer field.

## **W. Wallace McDowell Award**

### **Certificate/\$2,000**

Recent theoretical, design, educational, practical, or other tangible innovative contributions.

## **Taylor L. Booth Award**

### **Bronze Medal/\$5,000**

Contributions to computer science and engineering education.

## **Computer Science & Engineering Undergraduate Teaching Award**

### **Plaque/\$2,000**

Recognizes outstanding contributions to undergraduate education.

## **IEEE-CS/Software Engineering Institute Watts S. Humphrey Software Process Achievement Award**

*(Joint award by CS/SEI)*

### **Plaque/\$1,500**

Software professionals or teams responsible for an improvement to their organization's ability to create and evolve software-dependent systems.

**Deadline:** 15 October 2017

**Nomination Site:** [awards.computer.org](http://awards.computer.org)

**For more information visit:** [www.computer.org/awards](http://www.computer.org/awards)

