



ARTIFICIAL INTELLIGENCE FELLOWS PROGRAM

Insight Artificial Intelligence Fellows Program is an intensive, seven week professional training fellowship that bridges the gap between academic research or professional software engineering and a career in artificial intelligence. Based in Silicon Valley, New York City and Toronto, the program enables scientists, researchers, and engineers to learn the industry-specific skills needed to work in the growing field of artificial intelligence at leading companies.

Your Bridge to a Career in Artificial Intelligence

Are you an academic researcher or professional software engineer looking to transition into a career in artificial intelligence (AI)? Do you want a career that truly leverages your quantitative experience in a fast-growing, in-demand field that is making a positive impact in the world?

Top companies in many fields are hiring AI researchers and machine learning engineers skilled in deep learning, reinforcement learning and other cutting edge methods to help them create breakthrough products. People on these teams implement the latest machine learning research and integrate new models into production while ensuring scalability. While large scale datasets are becoming increasingly commonplace, and there exists an ever growing ecosystem of tools and technologies to deal with them, there is a severe shortage of talent to build and optimize these systems.


Insight Artificial Intelligence Fellows Program is a professional training fellowship that bridges the gap between academic research or professional software engineering and a career in artificial intelligence. This intensive full-time, seven week, in-person program enables experienced researchers, scientists, and software engineers to apply their existing skills to problems in AI, including, but not limited to computer vision, audio processing, and natural language processing. Fellows learn by building new models or scalable machine learning platforms in a self-directed, collaborative, and hands-on environment. Fellows are provided guidance and resources from industry and open source leaders in Silicon Valley, New York and Toronto. Additionally, Fellows gain access to a network of more than 750 Insight alumni, who are now working at over 200 companies. Immediately following the program, Fellows interview at leading companies in the San Francisco Bay Area, New York City, Greater Toronto Area, and receive continued guidance and support until they accept a full-time offer to join a top AI team.






Insight Artificial Intelligence in a Nutshell:

1. 7 week, full-time, professional Artificial Intelligence training fellowship in San Francisco CA, New York, NY, or Toronto, CAN.
2. Self-directed, project-based learning (no classes!) under the guidance of top industry AI experts.
3. A group of smart people who are excited about working on interesting problems while having a positive impact
4. Meet dozens of top companies, present your work to teams you're interested in, and interview immediately following the program.
5. Need-based scholarships available to help cover living expenses.



Why Artificial Intelligence?

The past few years have seen an unparalleled (and unexpected) jump in the effectiveness of advanced machine learning techniques, including deep neural networks, or deep learning. Three key factors for this leap are the advances in high performance computing, the development of distributed methods, and the availability of large labeled datasets. Most recently, these advances have made their way from the research labs to the applied engineering and product divisions of top companies and startups. The role of AI teams are significantly different from those of data engineering teams, which focus on data infrastructure and pipelines, and data science teams, which focus on asking the right questions and extracting meaningful information from data to make decisions or build data products. AI teams usually focus on building, optimizing, and scaling deep learning algorithms that emulate core human abilities such as vision, speech, language, decision making, and other complex tasks.

Tackling these problems in a scalable way requires developing new infrastructure and methods: companies need to effectively use more of their data for models, implementing algorithms that can be trained on distributed data and in a distributed manner. In addition, these models need to be deployed at scale, often requiring compressed models to work in cameras, phones, drones, and other mobile devices. There is a shortage of knowledgeable practitioners who can implement cutting edge deep learning algorithms (many of which are only written up in research papers) and scale them up to production level.

The projected impact of AI on the economy is unprecedented. McKinsey Institute estimates that these scaled algorithms “could have as much as \$5.2 trillion to \$6.7 trillion in economic impact annually by 2025.”¹ Many companies, from large brands to early-stage startups, now have a core products that revolves around AI (e.g. advanced natural language processing, computer vision) and are striving to keep up with or surpass the results coming from research labs, and put the latest AI algorithms to work on real world data. All of these companies are hiring researchers and engineers focused on building intelligent systems to augment their current practices. These AI professionals will be at the forefront of developing the next generation of products and solutions that will ultimately help millions of people in areas as diverse as healthcare, robotics and communications.

¹ McKinsey, Disruptive technologies: Advances that will transform life, business, and the global economy (2013)

Techniques, Trends, and Future Promise

“I think people need to understand that deep learning is making a lot of things, behind-the-scenes, much better.”²

- Geoff Hinton, Google Brain

Artificial Intelligence has a long and rich history, but while many of the tools and techniques have been around for decades (i.e. multi-layer perceptrons, convolutional neural networks, reinforcement learning), progress has been stalled due to the lack of sufficiently large scale datasets and the computational horsepower for efficient training. Recently, GPUs and distributed cloud computing systems have accelerated training, a number of massive well-annotated datasets have emerged to foster progress in many challenging visual problems (e.g. ImageNet, COCO, YFCC100M, Street View House Numbers, YouTube-8M, Open Images, and comma.ai’s self-driving car data), and new collaborative platforms have been engineered to encourage research in reinforcement learning (e.g. OpenAI Gym).

Due to these advances, there has been dramatic success in the development of intelligent systems. Here are a few of the recent tools and techniques that are reshaping the artificial intelligence landscape in the past couple years and are representative of the progress being made in the field:

- **Computer Vision:** Convolutional Neural Networks (CNNs) are deep neural networks specifically designed to take advantage of the structure of image input
 - In 2012, AlexNet built by a team at the University of Toronto became the first CNN to win the ImageNet challenge, an image recognition competition
 - From 2013 to 2015, UC Berkeley researchers created and improved R-CNN, the current state of the art method for object detection and recognition, by combining deep learning with region selection algorithms
 - In 2015, ImageNet competitors beat typical human performance with CNNs (e.g. Microsoft’s ResNet and Google’s GoogLeNet with batch normalization)
- **Reinforcement Learning:** Algorithms that learn action choices to maximize an objective
 - In 2006, researchers introduced a noisy cross-entropy method that outperformed existing algorithms in the game of Tetris
 - From 2013, Deep Q-learning was used to create agents that could successfully beat Atari games³
 - In 2015 and 2016, AlphaGo used deep neural networks combined with Monte Carlo Tree search and self play to defeat top Go professionals⁴

² Maclean’s, The meaning of AlphaGo, the AI program that beat a Go champ (2016)

³ Nature, Human-level control through deep reinforcement learning (2015)

⁴ Nature, Mastering the game of Go with deep neural networks and tree search (2016)

- **Advanced Natural Language Processing:** Neural networks to understand language
 - From 2012 onwards, Recurrent Neural Network (RNN) architectures redefined “state of the art” in the tasks of document summarization, language generation, and language to language translation (machine translation). Later LSTM models successfully generated language at the character level.
 - In 2013, Google created the word2vec algorithm, which provides a way to compare words based on their similarity in usage context. Further work in this area led to doc2vec (2014) and lda2vec (2016)
 - In 2015, researchers developed models to automatically generate captions of images by combining CNN and RNN architectures.
 - In 2016, the Google Neural Machine Translation system (GNMT) achieves state of the art results for machine translation
- **Unsupervised Learning:** Systems that learn the underlying structure of the data
 - In 2014 Generative Adversarial Networks (GANs)⁵ provided a new framework for synthesizing examples, leading to deep convolutional models (DCGANs)⁶
 - In 2014, advances in variational approaches led to tractable bayesian autoencoder models ⁷
 - The adaptation of RNN architectures in 2015 used localized attention to improve generated images⁸
- **Scalable Learning:** Distributed models and training of deep networks and AI agents
 - In 2012, Google developed a software framework and two algorithms for training distributed deep neural networks⁹
 - In 2015 and 2016, Google DeepMind developed distributed methods for training reinforcement learning agents ^{10,11}
 - In 2016, DeepMind used synthetic gradients, specialized units that model the true gradient, to decouple network layers and allow for parallelization
- **Emerging Trends**
 - Combining probabilistic programming with deep learning for uncertainty estimates, regularizing weights with new priors, and creating hybrid networks¹²
 - Leveraging GAN and Variational approaches for semi-supervised learning
 - Using memory networks to combine attention, memory, and reasoning for language understanding and developing Q&A systems

⁵ NIPS, Generative adversarial networks (2014)

⁶ Radford, et.al. Unsupervised representation learning with deep convolutional generative adversarial networks (2015)

⁷ ICLR, Auto-Encoding Variational Bayes (2014)

⁸ Gregor et al., DRAW: A Recurrent Neural Network For Image Generation (2015)

⁹ NIPS, Large Scale Distributed Deep Networks (2012)

¹⁰ ICML, Massively Parallel Methods for Deep Reinforcement Learning (2015)

¹¹ ICML, Asynchronous Methods for Deep Reinforcement Learning (2016)

¹² Thomas Wiecki, Bayesian Deep Learning (2016)

“AI is already changing our daily lives, almost entirely in ways that improve human health, safety, and productivity.” AI has the potential to “to make driving safer, help children learn, and extend and enhance people’s lives. In fact, beneficial AI applications in schools, homes, and hospitals are already growing at an accelerated pace.”

- Stanford’s first One Hundred Year Study on Artificial Intelligence

Here are some of the ways these techniques have now started to be deployed in industry:

- **Healthcare:** Researchers are improving the state of the art in cancer detection using neural networks on medical images. Companies are utilizing recurrent neural networks to uncover patterns in wearables data. Hospitals are poised to leverage advanced natural language processing methods to analyze large collections of medical records and control theoretic methods to optimize individualized medical trials and treatment plans.
- **Transportation:** Autonomous or near-autonomous vehicles are increasing the safety of motorists and helping to decrease global carbon emissions. Continued advances will dramatically alter the productivity shipping and trucking companies whose core business is the transport of goods. Companies are using reinforcement learning methods to grapple with the ever-increasing complexity of internal logistics and operations.
- **Robotics:** Sophisticated surgical robotics are advancing medical practice and the continual development of robotic prostheses are helping lives. Labs are developing techniques to help transition service robots from research to households, transforming elderly care. Increases in robotics in manufacturing can lead to more productivity.
- **Personalized AI Assistants:** Companies are leveraging advanced natural language processing methods and deploying chatbots to automate various aspects of their HR and customer service operations. There is an increasing number of startups that are rolling out interactive services that engage customers to provide individualized content recommendation.
- **Education:** Online educational platforms are providing individualized instruction schedules through reinforcement learning methods and instructors are using machine learning and natural language processing to scale up distributed educational efforts.
- **Satellite and UAV Imaging:** The decreasing cost of private satellites and drones is enabling the generation of massive amounts of imagery from which companies are deriving insights from and uncovering facts about the environment and societies around the globe.
- **Speech Recognition and Generation:** Twenty percent of Google search queries are voice queries¹³ and that figure is growing. Deep learning has transformed our voice to text capabilities. In other veins, research has helped advance machine translation systems as well as invent new systems to produce the most human-like text to speech systems on the market.

¹³ Sundar Pichai, Google I/O 2016 conference keynote (2016)

*“Deep learning produces real results and is at the root of a real industry that makes money today. The promises of it in the near future are very exciting... in areas like self-driving cars, medical imaging, personalized medicine, content filtering/ranking, etc.”*¹⁴

- Yann LeCun, Facebook AI Research

Deep learning and other techniques in AI, in addition to advancing the sectors above, will soon impact additional areas ripe with new challenges:

- **Video Understanding:** While there are a growing number of solutions for image analysis, including packages to provide detailed captions and descriptions, there is less work being done in understanding video content. As Geoff Hinton notes, “I think that the most exciting areas over the next five years will be really understanding videos and text. I will be disappointed if in five years time we do not have something that can watch a YouTube video and tell a story about what happened.”¹⁵
- **Security:** As individuals gain more skill in creating artificial agents that can act with malicious intent, it is ever more clear that systems will need to be implemented that can detect this anomalous behavior. Since the amount of activity is enormous, suspicious behavior is distributed in space and time, and the expertise need to analyze is unclear, there is a need to develop better unsupervised methods that can quickly learn activity patterns and surface relevant behavior
- **Public Safety and Security:** Stanford reports that “by 2030, [regional and national authorities] will rely heavily upon... improved cameras and drones for surveillance, algorithms to detect financial fraud, and predictive policing.” These will all rely heavily on the development of new methods and algorithms.
- **Internet of Things:** The ever-growing miniaturization of devices and incredible growth of sensor networks is providing a massive wealth of heterogenous signals which must be analyzed for patterns. Since it is unclear as to which features of this data are most relevant, methods that can automatically extract structure and provide predictive ability will be key.

¹⁴ Quora, Session with Yann LeCun (2016)


¹⁵ Reddit, AMA with Geoff Hinton (2014)

The Demand for Artificial Intelligence Professionals

“AI is the new electricity. Just as 100 years ago electricity transformed industry after industry, AI will now do the same.”

- Andrew Ng, Chief Scientist of Baidu

As the need for specialized expertise has grown and companies have continued to create new teams to address the optimization and scalability of advanced machine learning algorithms, the number of positions requiring an expertise with AI methods, including deep learning and machine learning engineering, has drastically increased.




Job Trends on [indeed.com](http://www.indeed.com/jobtrends) for deep learning, machine learning engineer ¹⁶

Top venture capital firms have been increasing their funding for AI startups over the last four years (see figure below).

¹⁶ [indeed.com](http://www.indeed.com/jobtrends) Job Trends search - <http://www.indeed.com/jobtrends>

“Private firms are investing heavily in AI — and in particular in an AI technique called deep learning — because of its promise to glean understanding from huge amounts of data.”

- Nature News, 2016



The demand for top AI professionals throughout industry, however, far exceeds the current supply. Top AI leaders have expressed their concern that “there are so many opportunities and so few people to work on them”¹⁷ — in other words: “everyone wants to hire ML engineers.”¹⁸ One reason for this gap is the combination of advanced software engineering and machine learning skills required for these roles. Another crucial issue is that, while many candidates have experience with foundational machine learning techniques, their knowledge of the latest research-grade techniques does not provide the requisite intuition for tackling real world problems in industry. Since AI methods are rapidly evolving and the data underlying business use cases are rare, scientists, and engineers rarely get a chance to calibrate to the status quo let alone make advances. Many companies are hesitant to hire candidates without this direct, product-focused machine learning experience, but many candidates with the necessary fundamentals cannot get that experience in their current roles.

This is where the Insight Artificial Intelligence Fellows Program comes in, and why top organizations are helping Insight develop the next generation of leading AI professionals.

¹⁷ Andrew Ng in Nature, AI talent grab sparks excitement and concern (2016)

¹⁸ Quora, Session with Yann LeCun (2016)

Who are the best AI Professionals?

The best AI professionals have a passion for solving complex machine learning challenges or engineering robust machine learning systems. We expect Insight AI Fellows to come from roughly two different category of backgrounds, although we're excited to consider an application from anyone, regardless of their background:

Quantitative Scientist: You have research experience at a postdoc, PhD, Masters or undergraduate level, have used advanced methods in machine learning to further your research and now want to pursue AI as a full-time career. Although your area of research may *not* be in computer science, you applied machine learning algorithms on messy data sets to further your research efforts and have an understanding of the underlying algorithms. You have a strong quantitative background and can read and interpret academic papers in machine learning and beyond. Likely fields of research can include physics, astrophysics, computational neuroscience, cognitive science, mathematics, computational biology, and any other quantitative, data-driven fields.

Software Engineer: You are a software engineer who has worked in machine learning either in an academic setting, in your current professional role or via building side projects. You are a professional software engineer or coming from an engineering or computer science academic program and you want to dive into the latest machine learning problems, demonstrate your abilities then join a top-tier AI team.

The top AI teams are employing a model where they are pairing researchers with engineers to leverage maximize for speed of experimentation and productionization. The Insight AI Fellows Program will be no different. By having both researchers and engineers in one room we will be able to collaboratively solve more sophisticated AI problems, making for an exciting learning environment for Fellows from both backgrounds.

To be successful, Fellows need the ability to synthesize current research, develop new models, iterate quickly in a software environment, develop large code bases, and enjoy building products. They also need to be able to put models into production, integrate novel machine learning algorithms into complex data pipelines, and ensure scalability. That said, you do not need to be strong in all the areas to apply. As part of the program you will learn quickly and fill in any gaps you have while leveraging the strength of experience you have already developed.

Insight Artificial Intelligence Program

In 2012, Insight developed a new model for education: we bring together very smart, hard-working, and enthusiastic scientists and engineers who have the fundamental skills, and enable them to make a transition into an advanced field by gaining hands-on experience with the tools used in industry and engaging with an extensive network of industry mentors.

Learning by Building

Since our Fellows enter the program with existing machine learning expertise and software engineering skills, there are no classes or lectures on these foundations. Instead, the main focus of the program is to learn by building models and scalable infrastructures using the same technologies and best practices as leaders in the AI community.

Projects will be suggested by leaders in the open-source machine learning community as well as startup companies funded by Y Combinator, Data Collective, Khosla Ventures and others (read more about the [Y Combinator + Insight partnership](#)). For more detailed examples of projects from previous Fellows in our data science, data engineering and health data programs, visit our [Insight blog](#).

Projects you can build while an Insight AI Fellow

Depending on their interests and company preferences, each Fellow build a project in one of the following areas:

Contribute to Open Source AI efforts:

- **Open Source Models:** Implement research-grade algorithm as a pre-trained model using an open source framework.
 - *Example:* implement a paper from a recent deep learning conference and incorporate it into the TensorFlow Model Zoo. Note: We'll have mentors from the Google TensorFlow team mentoring at Insight - you'll have the opportunity to potentially get featured on TensorFlow's website.
- **Open Source Methods:** Implement a production-quality contribution that adds functionality to an existing open source framework.
 - *Example:* implement and optimize a parallelized version of a sequential reinforcement learning algorithm.

Solve real world problems with AI:

- **Company-Specific Models:** Modify or extend current methods to address a real business or product application, which out-of-the box academic models are unable to tackle appropriately.
 - *Example:* modify an autoencoder network to subtract elements from an image
 - *Example:* alter a DC-GAN to fill in missing image elements
 - *Example:* demonstrate when and how to cut off layers of VGG to improve feature generation for a novel task.
 - *Example:* learn an optimal policy for truck routing to improve company operations
- **Data Manipulation:** Chart how manipulations to company data affect model performance and design methods to optimize performance.
 - *Examples:* develop methods to generate synthetic data to improve training or document how topological structure and training leads to differing internal representations.

Scale up AI methods:

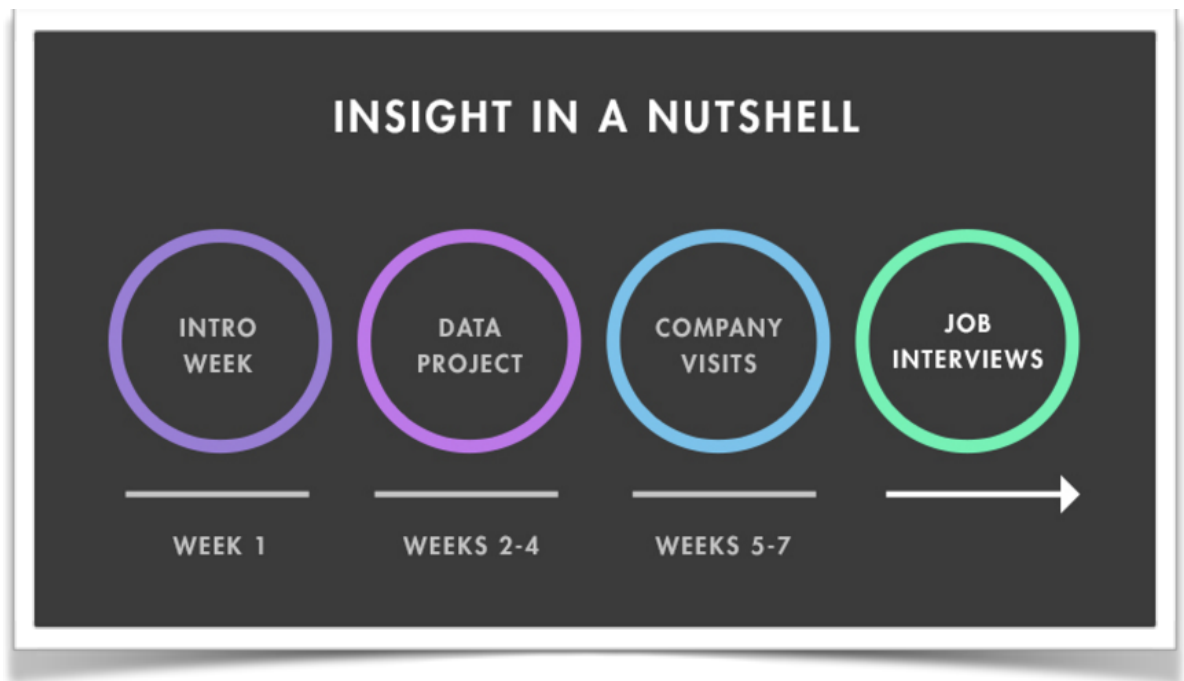
- **Learning at Scale:** Train advanced ML algorithms with distributed data and distributed models.
 - *Example:* using a distributed neural network (e.g. with synthetic gradients) and terabytes of distributed data, optimize the learning of a new neural network classifier by simultaneously and distributively learning the best hyperparameters.
- **Deployment at Scale:** Ingest and process data, apply advanced ML models to data, visualize with dashboards, and monitor the quality of results as data shifts over time. This emulates building an entire infrastructure for a startup scaling up their ML pipelines.
 - *Example:* Develop a pipeline that collects and analyzes millions of videos while allowing users to query those videos with objects to find them in frames throughout video collection.

Develop novel AI algorithms and approaches:

- **Novel Data:** Investigate new data sources that could benefit from recent advances in intelligent system design and determine some best practices for learning from that data.
 - *Example:* data from new IoT sensor networks and data from security companies are rich datasets that have yet to be fully explored.
- **Novel Methods:** Experiment with new types of network architecture that takes advantage of the structure of that data: 2D Convolutional layers are structured to better deal with visual information, the dilated causal convolutional layers of DeepMind's WaveNet are structured to better deal with audio data, and Matroid's V-CNNs are structured to better deal with volumetric data.
 - *Example:* develop a new type of network structure that could better deal with a specific data structure by manipulating it in a novel way.

Timeline

Insight is intended for scientists and engineers who are motivated to quickly build on their fundamental machine learning experience. Fellows gain meaningful experience with AI technologies and design patterns in just a few weeks. This high rate of learning mirrors the iterative nature of development in the tech industry.



Weekly Breakdown of the Program

- **Week 1: Intro** - Learn the AI ecosystem and begin building your data platform.
- **Week 2: MVP** - Build the Minimal Viable Project with the core functionality.
- **Week 3: Scaling** - Scale your project to handle more data and edge cases.
- **Week 4: Perfect and Practice** - Finish your project & practice presenting it.
- **Week 5-7: Present and Prep** - Present your AI project to the companies you are interested in, and begin preparing for the interviews that will follow.
- **Week 8+: Interview** - Interview at the companies that call you back from your presentations and sign an offer to join a top AI team in Silicon Valley, NYC or Toronto.

Who's involved?

The Insight Artificial Intelligence Fellows Program is a professional training fellowship that bridges the gap between academic research or software engineering and a career in artificial intelligence. With participation from leading technology companies, we are connecting top scientists and engineers with some of the most innovative companies in the world. Mentors for the Insight Artificial Intelligence Fellows Program are AI researchers and engineers at:



During the program, Insight AI Fellows will get to work on cutting edge applied AI problems provided by top AI startups funded by Y Combinator, Data Collective, Khosla Ventures and others:



Collaboration and Mentorship

Though each Fellow will build their own system based off their individual interests and abilities, the strength of Insight is rooted in a collaborative environment. Fellows accelerate their learning by working together to solve common problems and leveraging the diverse backgrounds of one another and the Insight network. Mentorship comes from the following sources:

- **Company Mentors** - AI teams at leading companies of various sectors and sizes that will share the problems they're currently solving. This helps you better understand the current challenges in the industry and decide which teams you want to join.
- **Industry and Open Source Leaders** - Pioneers at the forefront of Artificial Intelligence that help you learn the best practices and newest trends in the industry.
- **Your fellow Fellows** - Ambitious scientists, researchers, and engineers with common goals and a diverse set of skills that complement yours. Some have years of experience with databases and writing production quality code, while others studied sophisticated machine learning topics like natural language processing and computer vision. Learn the way you do in industry, by collaborating and working through challenges with your peers.
- **Insight Alumni** - Previous Fellows from our Data Science and Data Engineering programs have transitioned into AI roles at top companies. They provide individualized guidance and practice with interviews.
- **Insight Team** - Insight staff that offer continued guidance throughout the entire process. They point you to the right resources to help troubleshoot tough issues.

Insight Artificial Intelligence Program

The goal of Insight is to prepare the next generation of leading AI professionals. To do this, we have created a program that is explicitly designed to walk Fellows directly into fulfilling careers at companies on the leading edge of artificial intelligence. These companies have a very high bar for talent and are only looking for the best possible candidates in any position they hire for. This is why we are setting a very high standard for acceptance into the program and expect entry to be quite competitive.

While top-tier quantitative ability or software engineering experience is necessary, it is not sufficient to be an Insight Fellow. In addition to sheer smarts, we're looking for Fellows who are extremely curious people, highly motivated, love learning across a wide range of fields, enjoy collaborating with other smart, driven colleagues and are excited about the opportunity to make a positive impact in the world.

Responsibilities

As an Insight Fellow you're given the opportunity to learn from the best AI practitioners for seven weeks. The program is designed to remove as many obstacles as possible that stand between where you are now and becoming a rising AI professional. With these benefits, however, come a few responsibilities that you must be comfortable with before choosing to apply for the program.

- Actively and thoughtfully contribute to group activities and sessions during program hours Mon-Fri 10am-6pm.
- Take a leave of absence, if applicable, from your current responsibilities (e.g. current employment, research, or studies) in order to participate at Insight. Both the 7-week portion and the subsequent interview process require a full-time focus.
- Self-direct your learning and tackle a challenging Artificial Intelligence project during the program, while giving and receiving constructive feedback.
- Interview for full-time Artificial Intelligence positions with Insight mentor companies in your program's location upon completion of the program.
- Plan to keep coming into the office during the interview weeks (when not interviewing) to participate in interview prep sessions with other Fellows.
- Support future Fellows by providing mentorship and guidance once you become an AI professional in industry.

The guiding principle of Insight is: *Fellows first*. We strive to create an environment where you can learn and develop your career in artificial intelligence. All we ask in return is that you give it your all, be fully engaged in the process and help pass on your learning to the next batch of Fellows, helping to make the Insight community stronger as a result.

Benefits

The Insight program is designed to provide all the training, resources and connections you'll need to effectively transition to a career in artificial intelligence. Here are some of benefits of becoming an Insight Fellow

- Need-based scholarships are available to help cover living and travel expenses -- our goal is to make sure everyone with the right skills can participate in Insight, regardless of their financial situation.
- Desk space at the Insight offices in Silicon Valley, New York City or Toronto during the program, with full-time access to a library of relevant resources.

- Dedicated cloud computing resources to build and maintain your system for the duration of the program.
- Tips and assistance from our staff to help with your living arrangements for the duration of the program.
- Guidance and mentorship from industry professionals at every stage of the program and as you prepare for interviews.
- Mentorship from alumni Insight Fellows whose experience, at Insight and at their current data roles, make them an unparalleled resource to provide guidance and feedback.
- Personalized company matching. We help you figure out which companies are the best fit for you based on our experience and in-depth conversations with the hiring managers. We then help you arrange interviews during the final week of the program.
- Help navigating the negotiation of final employment terms once companies have made their employment offers to you.
- Perhaps most importantly: an unparalleled professional network of artificial intelligence professionals, data scientists, health data scientists, data engineers, friends and acquaintances. Through the program you will meet and get to know several dozen top AI professionals, who are Insight mentors and alumni, all of whom will be your industry peers. These professional contacts will be an invaluable source of knowledge, advice, career opportunities, and friendship in the years to come.

Insight Artificial Intelligence Program

Applications are currently open on [our website](#) for the next start dates in each location of Insight Artificial Intelligence Fellows Program. We expect high demand, so we encourage you to apply early.

If you have any questions please email us at info@insightdata.ai

Or to apply, please visit: <http://insightdata.ai/apply>