#1 TOP AI BLOG
TOP 50 BLOGS

aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning –
a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing
– Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

AITrends.com

aitrends
The Business and Technology of Enterprise AI

Trends in The Artificial Intelligence Era:

EVERYTHING YOU NEED TO KNOW

Sponsored by deepsense.ai BIG DATA SCIENCE

Produced by Cambridge Innovation Institute

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning –
a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing
– Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

deepsense.ai
BIG DATA SCIENCE

Trends in the Artificial Intelligence Era: Everything you need to know

Artificial Intelligence (AI) is a new and powerful tool that is used to solve a wide range of problems in creative and often unexpected ways, and it is developing fast. So fast, in fact, that it has become indispensable for any business with designs on becoming an industry leader.

In the upcoming years, managers will have to focus on addressing their customers' needs in innovative ways. Those needs will be both well-known ones and others that have yet to appear. Their innovativeness will help organizations build strong brands and secure long-term growth. Working out smart and hard-to-copy methods of delivering the highest standards in service will drive the next decade of enterprise innovation.

In this overview, experts from deepsense.ai present insights from four areas to help managers meet their goals. eCommerce, retail, banking, marketing, entertainment, healthcare, and medical diagnostics are among the industries that stand to benefit the most from the trend analysis. However, these cases can provide inspiration for people from throughout the modern economic and professional landscape.

This eBook presents effective segmentation systems for better communication with your target audience, how machine learning boosts medical diagnostics to support trained clinicians, as well as an overview of reinforcement learning and how you can start benefiting from AI's nascent brainchild.

The book also details one of the most dynamic areas of AI, natural language processing (NLP), which is providing enterprise with an array of opportunities. Read up on how you can incorporate it into your business and what challenges you'll need to overcome.

We encourage you to follow us on our blog for more insights and to share your questions by email at contact@deepsense.ai.

Here's wishing you a good read!

The deepsense.ai team

ABOUT AI TRENDS | AITrends.com

AI Trends is the leading industry media channel focused on the business and technology of enterprise AI. It is designed for business executives wishing to keep track of the major industry business trends, technologies and solutions that can help them keep in front of the fast-moving world of AI and gain competitive advantage.

CONTACTS:

PRESIDENT
Phillips Kuhl

PUBLISHER
Lisa Scimemi
lscimemi@cambridgeinnovationinstitute.com

EXECUTIVE EDITOR
Eliot Weinman
ew@aiworld.com

EDITOR
John Desmond
jd@aiworld.com

VICE PRESIDENT, BUSINESS DEVELOPMENT
Angela Parsons
aparsons@cambridgeinnovationinstitute.com

MARKETING ASSOCIATE
Tom Conley
tconley@cambridgeinnovationinstitute.com

MARKETING ASSISTANT
Brooke Trefry
btrefry@cambridgeinnovationinstitute.com

MARKETING ASSISTANT
Renata Katz
rkatz@cambridgeinnovationinstitute.com

CONTACT INFORMATION
info@aitrends.com

250 First Avenue, Suite 300
Needham, MA 02494

CONNECT WITH US:

#1 TOP AI BLOG
TOP 50 BLOGS

aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning – a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing – Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

AI Customer Targeting Levels

From Customer Segmentation to Recommendation Systems

Every time you watch a film Netflix has suggested or buy a "similar product" on Amazon, it is a personalized recommendation. Can you make such recommendations work for your business as well?

Currently there are four levels of advancement in customer targeting, from no segmentation at all to advanced recommendation systems.

1. No segmentation at all
targeting all potential customers the same way

2. Manual segmentation
the most intuitive technique, the segmentation being done by human analysts

3. Automated segmentation
using machine learning to segment datasets and look for hidden patterns

4. Recommendation systems
instead of building a limited number of segments, these systems build an individual representation of each customer and product

Each of the four approaches has unique benefits.

Learn More

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning – a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing
– Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

1. NO SEGMENTATION

In the age of the Internet, treating all customers as a homogenous group will nick your popularity. This is why 83% of companies use at least a basic form of segmentation in their daily business. On the other hand, 43% of marketers don't send targeted emails.

Sometimes it is just not necessary to do so. If the business is a highly specialized or niche one or involves companies with few customers, further segmentation would not provide a significant return.

- When an online bookstore sells only legal publications, there is usually no need to segment customers, as only lawyers or professionals may need such books.

The need for segmentation grows alongside the scale of the business, as even with the most narrow segment customers are not a homogenous group and their needs may differ.

BENEFITS:

- Simple and cheap
- Effective at the beginning, when there are few customers
- No costs to maintain or implement

DRAWBACKS:

- Ineffective at larger scale
- Inflexible
- Lost opportunities resulting in low effectiveness

WHO CAN BENEFIT:

- Small companies with few customers or companies with a narrow target group and high specialization

"THE NEED FOR SEGMENTATION GROWS ALONGSIDE THE SCALE OF THE BUSINESS, AS EVEN WITH THE MOST NARROW SEGMENT CUSTOMERS ARE NOT A HOMOGENOUS GROUP AND THEIR NEEDS MAY DIFFER."

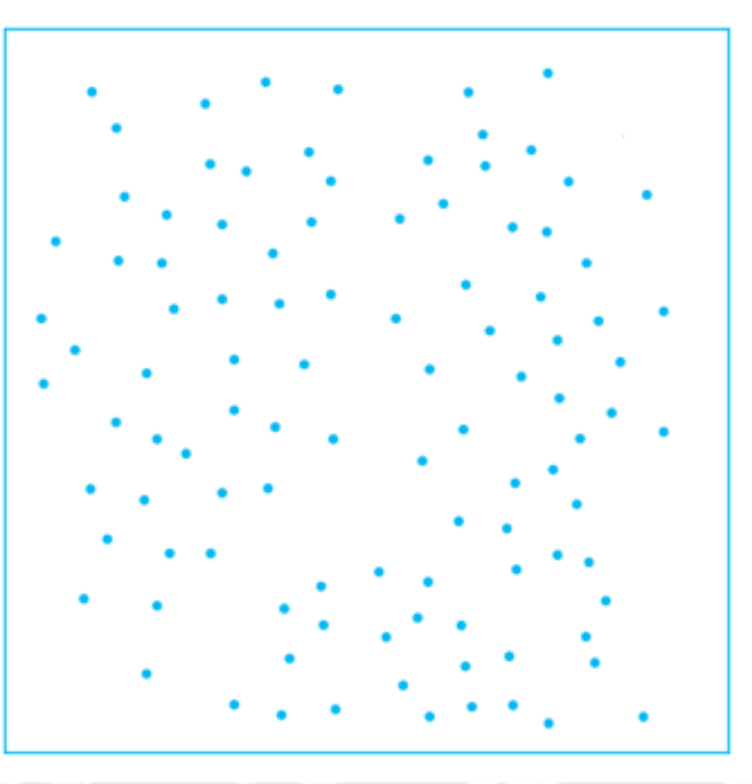

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO
SPONSOR'S LETTER
FEATURED ARTICLES
AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems
Reinforcement Learning –
a Moonshot or Today's Most Underhyped Technology
Machine Learning Advancing Medical Imaging and Analysis
The State of Natural Language Processing
– Giant Prospects, Great Challenges
aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

2. MANUAL CUSTOMER SEGMENTATION

Human analysts using tools of varying degrees of complication can handle segmentation manually: from Excel sheets to Tableau and advanced Business Intelligence tools. The analysts usually look at intuitive segments, for example demographic divisions (age, gender) and other criteria including geography, income, total purchase value or other factors.

- An online bookstore selling popular literature could segment readers into three groups: youth, women and men, each of whom obviously have their own preferences.

Even with all the benefits that attend building groups of customers, doing it manually presents significant challenges:

- Analysts processing the data may be biased. Teenage boys are stereotyped as computer gamers, but mature women play computer games more than young boys.
- With the dynamic nature of the market, every manual segmentation quickly becomes outdated. Then the work needs to be redone – again manually.
- The number of groups and segments researchers can create, validate and maintain is limited.
- Manual segmentation is not scalable

Even considering all the challenges the manual customer segmentation comes with, it is still a powerful efficiency boosting tool. For many companies, including most small businesses, manual segmentation is just enough.

BENEFITS:

- Intuitive and simple
- Greater efficiency than no segmentation at all
- Transparent and easy to understand

DRAWBACKS:

- Not scalable
- Fairly inflexible
- Requires constant maintenance, updates and supervision by human analysts

WHO CAN BENEFIT:

- Small, middle and sometimes large companies with easily segmented customer groups, including companies selling products tailored for demographics or using other straightforward criteria to target customers

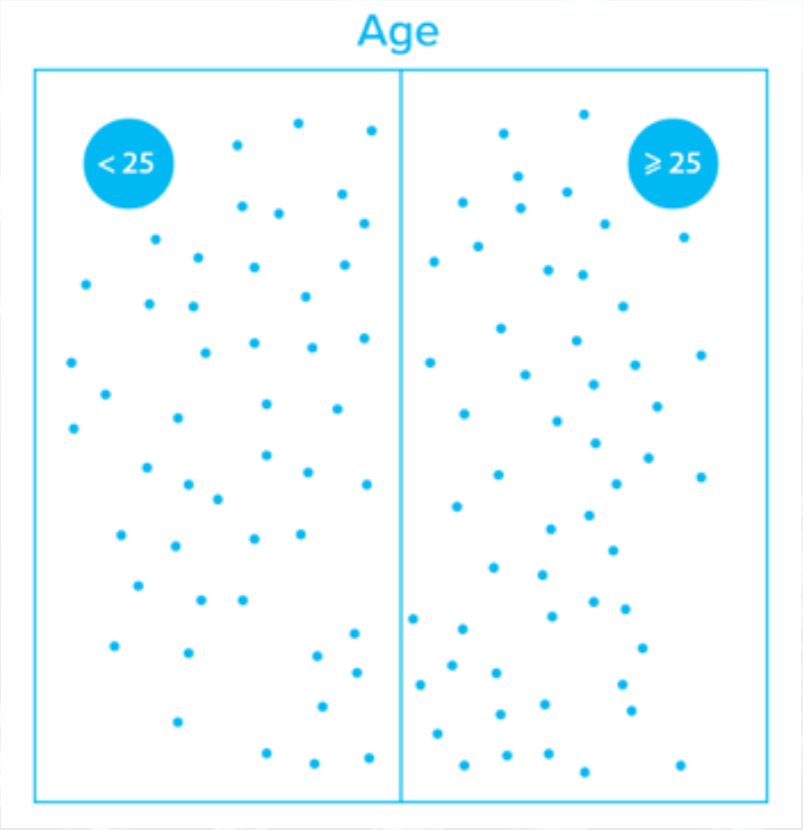

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO
SPONSOR'S LETTER
FEATURED ARTICLES
AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems
Reinforcement Learning –
a Moonshot or Today's Most Underhyped Technology
Machine Learning Advancing Medical Imaging and Analysis
The State of Natural Language Processing
– Giant Prospects, Great Challenges
aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

3. AUTOMATIC SEGMENTATION DONE WITH MACHINE LEARNING

Machine learning can be used to predict behaviour such as affinity for a given product or churn probability. However, this approach becomes slightly more challenging if you want to cluster similar customers, when there is no "ground truth".

K-means and hierarchical aggregation are currently the most widely used algorithms to cluster datasets without human supervision. Each point (customers in the dataset) is assigned to a class. Leaving behind the limited perception of the human researcher, including hidden biases and presuppositions, these algorithms can spot the most obscure and surprising and least obvious clusters within the dataset.

- In k-means, the number of clusters is fixed and the algorithm finds neighbouring points accordingly, as explained in this visualization.
- Hierarchical aggregation is an array of techniques that connect neighbouring points step-by-step, forming a dendrogram. While the algorithm provides more flexibility (we can set the number of clusters post factum), it requires stricter supervision and the clustering may be less stable.

One double take-inducing example comes from a man who used clustering to find the perfect woman for him from OkCupid's database. Such inspiring applications aside, using algorithms to segment datasets will present a number of challenges:

- Clustering is done on all data, both the useful and the irrelevant (e.g. hair color and complexion may be useful for shampoos but useless for taste in films)
- There is a fixed number of clusters. Although there are heuristics designed to tackle the problem, it remains an arbitrary choice. Clustering puts everyone in a distinct group, but there are surely more than 50 shades of grey between black and white.
  - Establishing a sharp line between readers of fiction and non-fiction books is every bit as hard as distinguishing between pure high-fantasy lovers and hard science fiction readers.
- Usually there is an actionable interpretation only for some groups..

It is therefore sometimes better to treat every customer individually instead of building groups and trying to find which group is the best fit. That's the bedrock of recommendation systems.

BENEFITS:

- Finds hidden clusters within the dataset
- Automated and therefore has no presuppositions
- Easier to scale and maintain productivity

DRAWBACKS:

- Requires human supervision, further interpretation of the segments, may produce segments that make no sense
- Requires maintenance and updates

WHO CAN BENEFIT:

- Larger organizations with too much data to handle manually

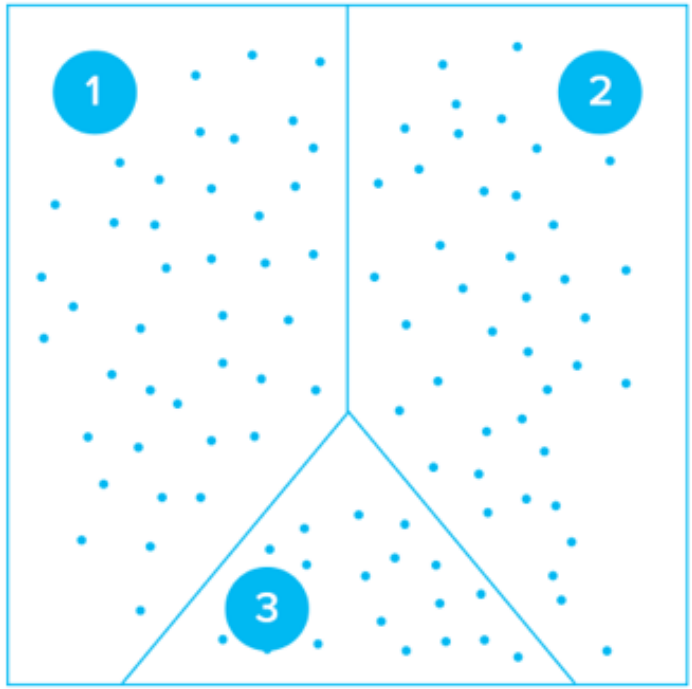

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO
SPONSOR'S LETTER
FEATURED ARTICLES
AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems
Reinforcement Learning –
a Moonshot or Today's Most Underhyped Technology
Machine Learning Advancing Medical Imaging and Analysis
The State of Natural Language Processing
– Giant Prospects, Great Challenges
aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

4. RECOMMENDATION SYSTEMS

Instead of seeking groups within a dataset, recommendation engines provide the customer with a representation in the form of a multidimensional vector (much like word2vec for discovering word similarities and analogies). It shows how the customer is perceived in terms of inferred (not chosen!) factors:

- How much do our book readers like fiction or non-fiction books?
- What's their attitude to fantasy and science fiction novels?
- How do they feel about romantic, dramatic and action-packed plots?
- What is the writer's political affiliation?

Modern recommendation systems, such as Factorization Machines, are able to leverage both official data (the author, genre, date of publication) and less obvious information inferred from buying patterns (are there any fictional monsters in the plot?, are there any supernatural horror elements?).

- If the reader likes political fiction, is slightly interested in science fiction and loves dramatic stories, he would probably be keen to read Margaret Atwood's "The Handmaid's Tale",

Assigning vectors to both customers and products allows the company to build narrow segments within segments and launch precisely targeted marketing campaigns.

- Instead of "fiction weekend", the online bookstore targeted narrow segments with the books individuals in a given group prefer. The group of fantasy novel readers got information about "high fantasy weekend" and selected 10 books that may be the most interesting for them.

We can take things even a step further, showcasing not only similar products, but analogous ones too – for example, a maternity version of a dress a client likes.

The recommendation systems also work with datasets that are not fully covered. No one reads all the books that are published. But a dozen books may be enough to predict what a person will like.

Using vector representation makes it is easier to validate if the system is working properly – every purchase is a feedback and may be used to readjust the model.

BENEFITS:

- Every customer is treated individually
- Constantly updated and evaluated
- Scalable

DRAWBACKS:

- Requires vast amounts of data to work properly
- The technology is complicated and requires skilled data scientists

WHO CAN BENEFIT:

- Large organizations with data-oriented culture, which process great amounts of data and are able to leverage the system.

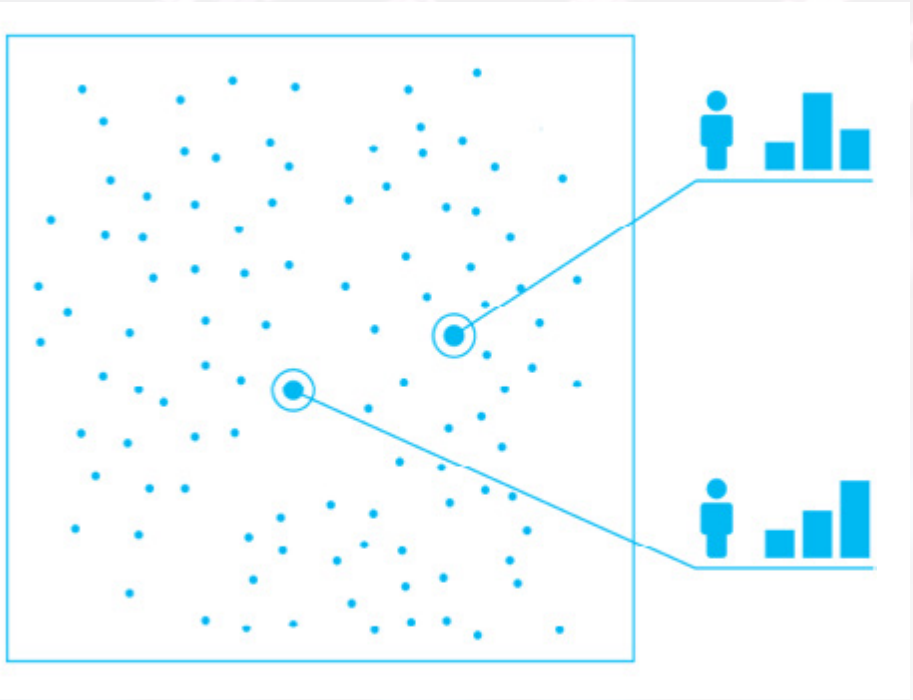

#1 TOP AI BLOG
TOP 50 BLOGS
aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning – a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing
– Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

Reinforcement Learning – a Moonshot or Today's Most Underhyped Technology

Reinforcement learning is gaining attention as the "next step in AI", but there are very few business use cases of this technology. So is reinforcement learning a moonshot or an underhyped game-changer?

Artificial intelligence is one of the most dynamic fields of research and development. It allows companies to solve new classes of problems and effectively tackle challenges where creativity and flexibility are required.

Unlike "machine learning" and "deep learning", "reinforcement learning" is not a commonly used buzzword and requires some explanation.

WHAT IS REINFORCEMENT LEARNING?

Data is fundamental to all forms of machine learning. Whatever model is being used, it needs to get and process data to be able to perform a task. In traditional machine learning and deep learning, the data usually come in the form of a closed dataset consisting of more or less structured and homogenous information. They may be a set of images, a sorted sheet or a customer database.

In traditional machine learning, the machines are given the ability to progressively improve their performance on a given task. This can be done with or without supervision:

- **Supervised machine learning –** when the ground truth labels for training input are at least partially available to the system. Using these data, the machine is able to learn the desired mapping and make predictions about events still to come. A good example is predicting seismic events in coal mines.
- **Unsupervised machine learning –** if the learning algorithm is provided only data without any labels, and tasked with finding the hidden structure and relationships within that data. This can be used to better understand and visualize data, or detect anomalies.

A great example of supervised machine learning (and deep learning) is the whale recognition model deepsense.ai developed for the US National Oceanic and Atmospheric Administration, an American scientific agency protecting endangered ocean species. The challenge was to help track the shrinking population of North Atlantic right whales by recognizing and distinguishing particular whales from aerial photographs. The model used the same mechanism as facial recognition, analyzing the photograph to look for specific patterns which it learned from the data provided.

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning –
a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing
– Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

THE REWARD HUNTER

In reinforcement learning, data are generated when an agent explores the environment. But the agent's behavior is shaped by a system of rewards and punishments. Building the artificial intelligence that controls an autonomous car is a great example – the agent is going to control the machine in a rapidly changing, unpredictable environment.

The system of rewards is a key to success – the agent gets points for safe driving and sticking to the rules. Crashing, running over pedestrians and speeding all result in penalties. These rewards serve as feedback the algorithm uses to get better at performing the task. They are somehow analogous to ground truth labels in traditional supervised learning, yet this time may be served in a much less friendly way e.g., delayed, sparse, or short-term. Reinforcement learning agents usually are trained in a simulated environment. After all, crashing a few dozen cars to teach an agent how to break would be prohibitively expensive.

So when is reinforcement learning actually useful and more effective than other AI techniques?

WHEN TO USE REINFORCEMENT LEARNING

Reinforcement learning solves a particular class of problems, especially ones that explore the environment or possible actions to perform. It deals with problems that traditional machine learning cannot handle:

- **Open –** challenges that allow multiple solutions. While traditional machine learning is effective in solving closed problems like image, speech or text recognition, reinforcement learning solves problems in the open world, where there are no correct answers.
  - **Example –** there is no "proper route" a car should take from point A to point B. The algorithm may look for the fastest, the safest or the most visually pleasing way to go, but this is not a 0-1 problem.
- **Possible to simulate properly –** reinforcement learning explores an environment, and builds the agent's skills by doing so. This means that using a proper simulator is crucial to eventual success.
  - **Example –** In the "Learning to run" project, muscles and bones were simulated thanks to the OpenSim simulation environment used in medicine. So crucial was that environment that the project would have made no sense without it.

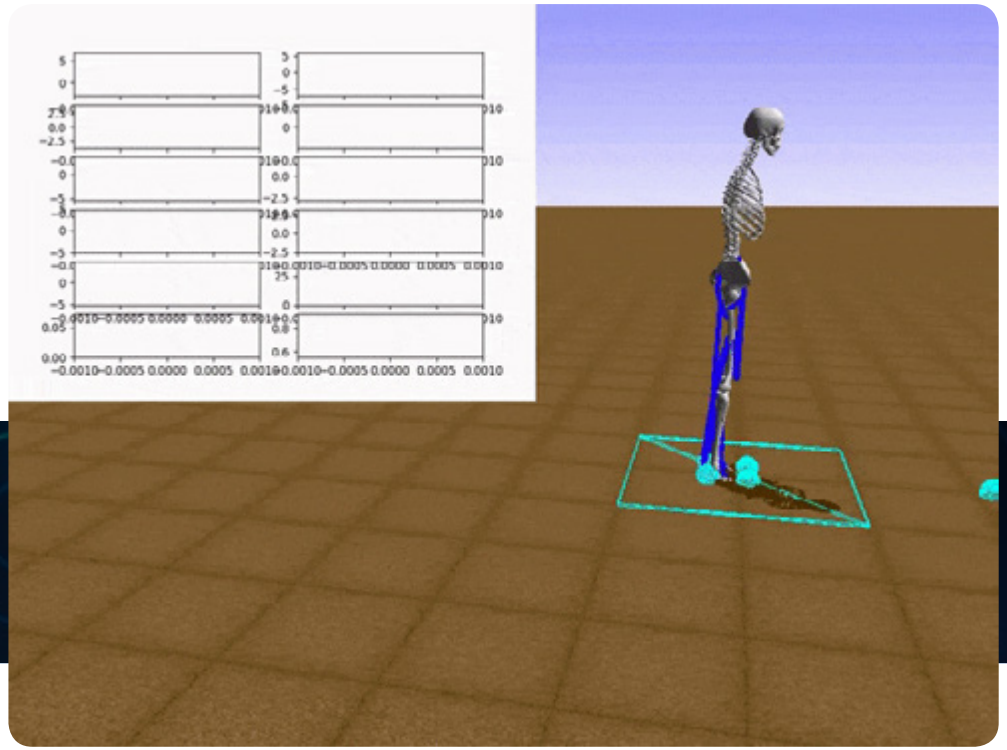

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO
SPONSOR'S LETTER
FEATURED ARTICLES
AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems
Reinforcement Learning – a Moonshot or Today's Most Underhyped Technology
Machine Learning Advancing Medical Imaging and Analysis
The State of Natural Language Processing
– Giant Prospects, Great Challenges
aitrends The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

- **Highly unpredictable –** when it comes to image recognition or text processing, the level of unpredictability is limited. The neural network is trained to recognize a particular pattern and it classifies the input into one of the categories. On the other hand, there are countless situations on the road or in other open environments that are impossible to predict, so hard-coding the conditions is virtually impossible. Worse, each action of the algorithm may inevitably change the environment, making any anticipation way harder.
  - **Example –** There is no way to predict what may happen on a road. Thus, the algorithm needs to be flexible and adaptable to react properly to unexpected conditions.
- **Overwhelmingly complex –** building software to control robotic arms or autonomous cars, while possible, is a complex undertaking that will be fraught with problems. By forcing the machine to develop desired patterns of behavior by itself, the developers save time and effort.
  - **Example –** A robot's behavior can be hard coded, but the process is complicated and full of potential mistakes. Forcing the machine to learn by itself makes the process less troublesome, as the machine is given a goal to achieve rather than instructions for how to do it.
- **Creative –** Traditionally, machines have approached challenges with neither creativity nor flexibility. That's why projects involving playing chess or go are so "hot" now – the machine needs to adapt and react to its opponent's strategy. Solving problems by developing a strategy is the highest level of creativity machines are able to achieve.
  - **Example –** flexibility is sometimes about breaking the rules. Reshaping the system of rewards may be used to build the machine that is able to respond appropriately to even the most unexpected situations. For example, in emergency situations speeding may be acceptable, but running over a pedestrian will never be.
- **Sequential –** many problems can be solved by a series of actions that need to be performed one by one in a particular order. Reinforcement learning handles such tasks with ease, while automating the sequences still requires a lot of work from programmers.
  - **Example –** In one particular case, a robotic arm learned to take a can of coke from a refrigerator. It needed to open the door, find the can on a shelf, grab it and close the door. In this scenario, there was a sequence of actions to be performed.

Building a flexible and creative machine is an achievement itself and reinforcement learning is undoubtedly perceived as an important next step in the development of machine learning. Of course, being "the next step" doesn't imply that other techniques are "of a lower generation". They just work for different classes of problems.

The agents deepsense.ai has trained to walk like a human or play Space Invaders go a long way towards illustrating the terms above. There are no particular ways to solve these problems, the world involved is open and the challenge is too complex to hard-code a solution.

"BUILDING A FLEXIBLE AND CREATIVE MACHINE IS AN ACHIEVEMENT ITSELF AND REINFORCEMENT LEARNING IS UNDOUBTEDLY PERCEIVED AS AN IMPORTANT NEXT STEP IN THE DEVELOPMENT OF MACHINE LEARNING. "

#1 TOP AI BLOG
TOP 50 BLOGS
aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning –
a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing
– Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

THE REINFORCEMENT LEARNING SUPERPOWERS

By combining interactions with the environment and machine learning, reinforcement learning provides interesting ways of development:

- By combining reinforcement learning and traditional supervised machine learning, it is possible to augment machine learning with expert knowledge. The neural network is fed with expert input, and reinforcement learning provides the flexibility to use the knowledge in the new context. What's more, the expert needn't be a neural network – it may be a human or other, non-AI-powered machine. Teaching an agent to play Montezuma's Revenge was one of the first examples of these two types of learning working together.
- It is possible to use reinforcement learning to build a neural network that mimics the surrounding world. It may be used to provide a cheaper training environment for the agent – instead of running the full simulation environment, the data science team needs to feed one neural network with the outputs of a second one. The details of the experiment and how it can be used may be found on deepsense.ai's blog.
- Reinforcement learning agents are able to outperform humans in many games, from chess to Atari classics. Thus, reinforcement learning is suitable for every situation where the goal is to win a game and the ability to react quickly to rapid changes is needed.

THE DOWNSIDES OF REINFORCEMENT LEARNING

There are classes of problems where the reinforcement learning agent will not be as effective as traditional machine learning techniques. Reinforcement learning is also expensive and challenging to use, so applying it to problems that can be solved effectively with other means makes little sense. What's more, a static environment cannot leverage the adaptability reinforcement learning agents offer. Further, training the reinforcement learning agent requires gobs of computational power for both the agent and the environment to train it. Lastly, finding specialists who know how to train reinforcement learning agents is no easy task as the discipline is only now gaining steam.

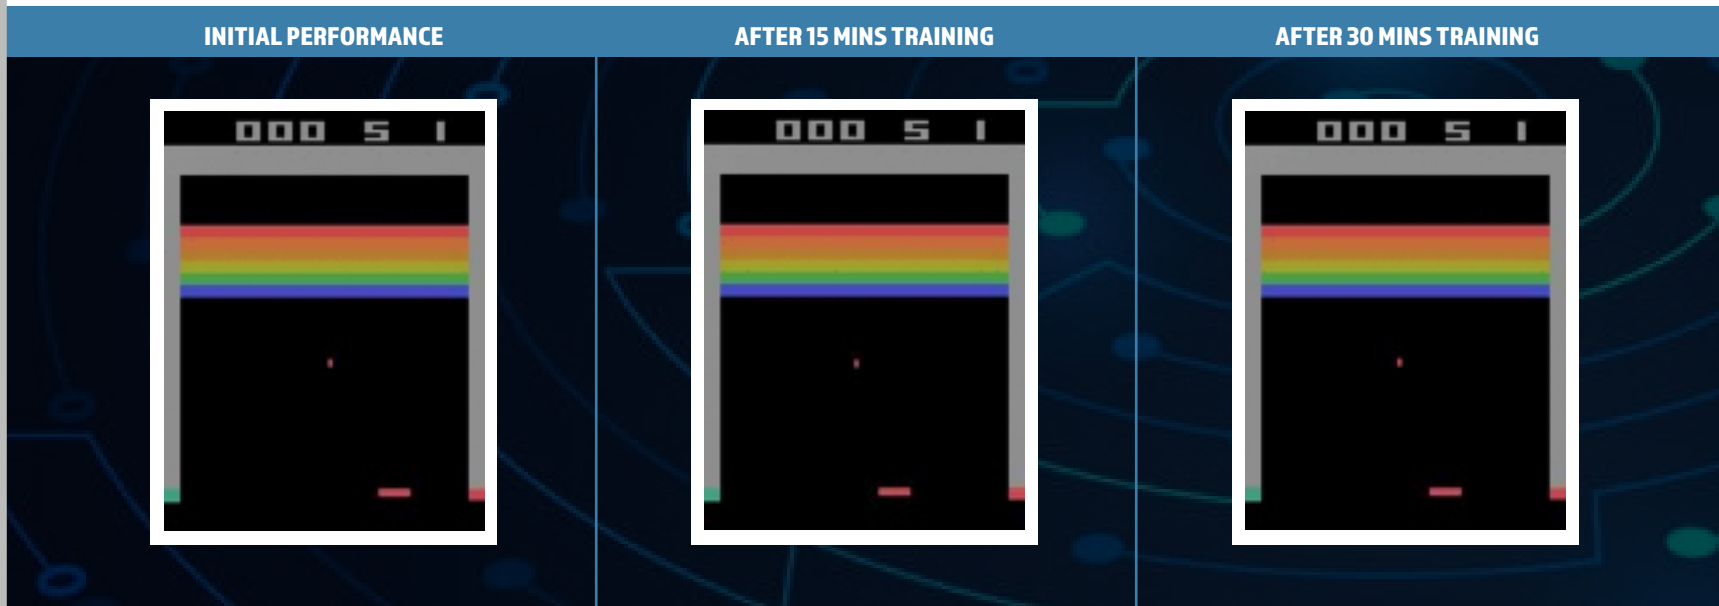

#1 TOP AI BLOG TOP 50 BLOGS

aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning – a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing
– Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

Machine Learning Advancing Medical Imaging and Analysis

By Pawel Godula,
Director of Customer Analytics,
deepsense.ai

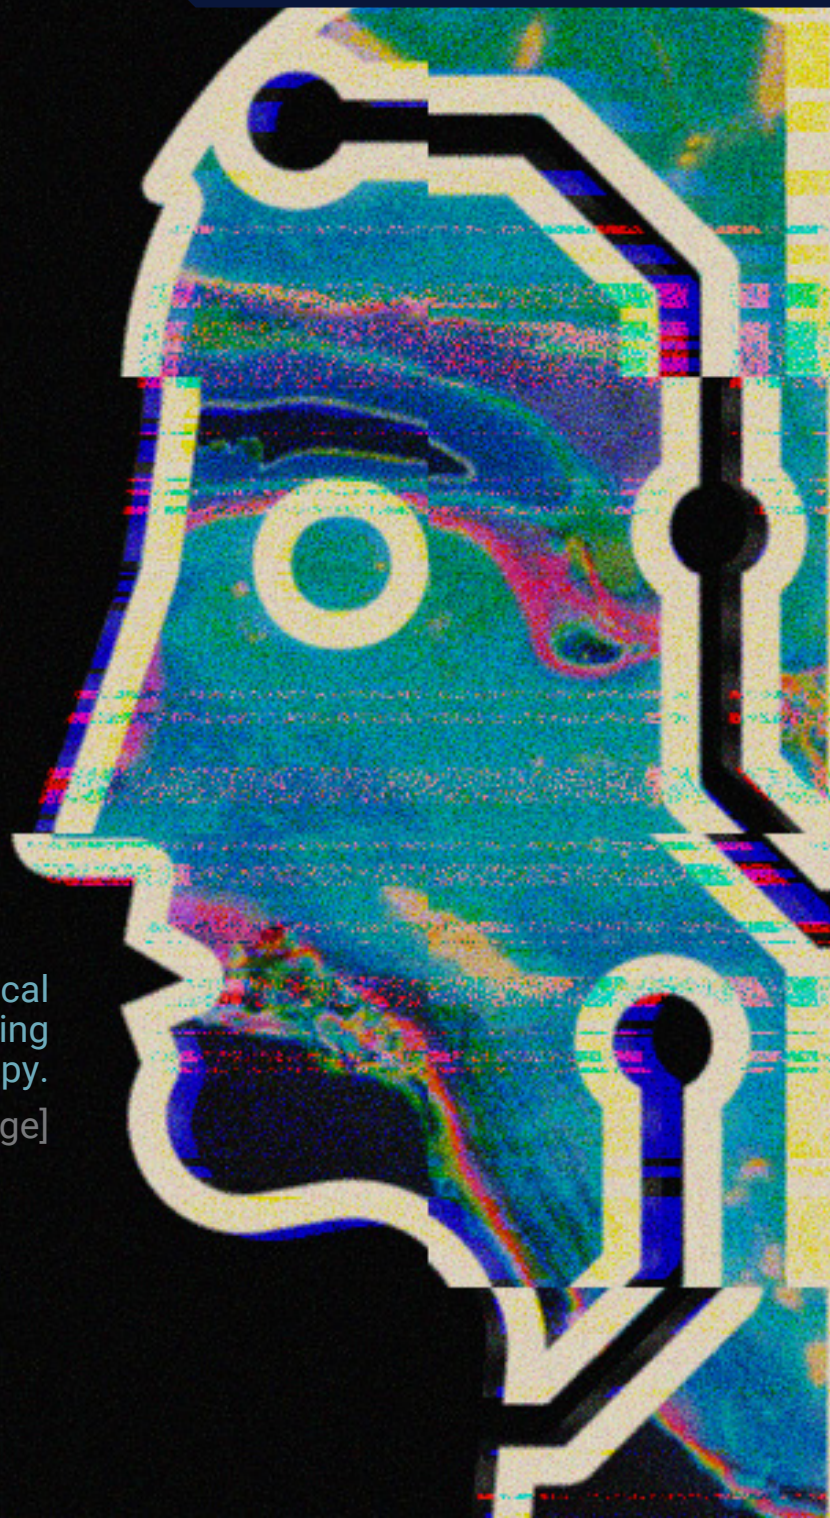

Machine learning is useful in many medical disciplines that rely heavily on imaging, including radiology, oncology and radiation therapy.

[Continued on Next Page]

#1 TOP AI BLOG TOP 50 BLOGS

aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning – a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing – Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

According to IBM estimations, images currently account for up to 90% of all medical data. Due to recent advancements, image recognition, especially with transfer learning done with networks pre-tuned on an ImageNet dataset, provides interesting possibilities to support medical procedures and treatment.

AI startups are being acquired at an increasing rate, while the value of AI healthcare-related equipment is also growing rapidly. As Accenture estimates show, the market is set to register an astonishing compound annual growth rate (CAGR) of 40% through 2021. Meanwhile, the market value of AI in healthcare is projected to skyrocket from $600M in 2014 to $6.6B in 2021.

Automated image diagnosis in healthcare is estimated to bring in up to $3B. Unlike many improvements that have been made in healthcare, AI promises both enhancements and savings. It can tackle common image-related challenges and automate heavy data-reliant techniques, which are usually both time-consuming and expensive.

DATA LABELLING AND A SKILL GAP

One of the most significant challenges in image recognition is the labor-intensive data labellingthat precedes the building of any new image recognition model. See our recent blog post concerning transfer learning.

Fortunately, some medical image data is spared. Radiological descriptions, for example, are standardized, applying a golden format to apply machine learning algorithms due to the labeling of data and enforcing order within the dataset. A challenge in modern radiology is to use machine learning to automatically interpret medical images and describe what they show. However, as the history of ImageNet shows, providing the properly labeled dataset is the first step in building modern image recognition solutions.

According to the American Journal of Roentgenology, if machine learning is to be applied successfully in radiology, radiologists will have to extend their knowledge of statistics and data science, including common algorithms, supervised and unsupervised techniques and statistical pitfalls, to supervise and correctly interpret ML-derived results. To address the skills gap among radiologists, companies that can handle the data science side of the equation, including teaching it, will be among the best solutions.

THE RISE OF RADIOGENOMICS

Combining different types of imaging data with genetic data could bring about better diagnostics and therapy – and potentially be used to uncover the biology of cancer. The new discipline of radiogenomics connects images with gene expression patterns and methods to map modalities. The paper entitled decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach describes an example of the process.

Interestingly, both image recognition (IR) and natural language processing (NLP) techniques can be used to analyze genetic data. Image recognition can be applied when the genomic data presents a one-dimensional picture consisting of colors representing each gene. The algorithms used are similar to any other image recognition approach. As machine learning models consider size irrelevant, among other factors, models may shape up to be similar as described in our recent blog post. NLP is used when the genes are represented by letters. While it is inferior to image recognition in looking for patterns and general analysis, NLP is better at seeing "the bigger picture" and looking for longer patterns present in larger sequences of genes.

MACHINE LEARNING IN PRECISION RADIATION ONCOLOGY

Radiogenomics is also an emerging discipline in precision radiation oncology. Machine learning approaches can be used to study the impact of genomic variations on the sensitivity of normal and tumor tissue to radiation.

#1 TOP AI BLOG TOP 50 BLOGS

aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning – a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing – Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

Radiation oncology is particularly well suited for applying machine learning approaches due to the enormous amount of standardized data gathered in time series. Radiotherapy involves several stages encompassing the entire oncological treatment:

- patient assessment,
- simulation, planning,
- quality assurance,
- treatment delivery,
- follow-up

All these stages can be supported and enhanced with machine learning. Tumors may have subregions of different biology, genetics and response to treatment. Thus, it is crucial to find spaces on images that need to be radiated with lower doses to make the therapy more precise and less toxic.

BUILDING MEDICAL IMAGE DATABASES – A CHALLENGE TO OVERCOME

Having access to proper datasets is a challenge to be tackled in medical image analysis. To gain insight into the mechanism and biology of a disease, and to build diagnostic and therapeutic strategy with machine learning, datasets including imaging data and related genetic data are needed.

According to Advances in Radiation Oncology, there are numerous databases and datasets containing healthcare data, yet they are not interconnected. Gaining high quality datasets containing medical data is quite a challenge and there are very few such datasets available. A collection containing images from 89 non-small cell lung cancer (NSCLC) patients that were treated with surgery is one of very few examples. For those patients, pretreatment CT scans, gene expression, and clinical data are available.

Efforts to build proper databases to support analysis of imaging data are being made. ePAD is a freely available quantitative imaging informatics platform, developed at Stanford Medicine Radiology Department. Thanks to its plug-in architecture, ePAD can be used to support a wide range of imaging-based projects. Also, TCIA is a service that hosts a large number of publicly available of medical images of cancer. The data are organized as collections including:

- patients related by a common disease,
- image modality (MRI, CT, etc.),
- research focus

Advances have already been made in histological image analysis and its clinical interpretation. deepsense.ai work has proved that it is possible to accurately analyze and interpret the medical images in diabetic retinopathy diagnosis. deepsense.ai built its model in cooperation with California Healthcare Foundation and a dataset consisting of 35,000 images provided by EyePACS.

Using this technique is more common. A machine learning approach reveals latent vascular phenotypes predictive of renal cancer outcome based on analysis of vessels in histological images. Vascular phenotype is related to biology of cancer. Forming new vessels is kind of a predictor–biomarker for potential of cancer development.

HAVING ACCESS TO PROPER DATASETS IS A CHALLENGE TO BE TACKLED IN MEDICAL IMAGE ANALYSIS.

#1 TOP AI BLOG TOP 50 BLOGS

aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning – a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing – Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

MODERN EQUIPMENT

The effectiveness of machine learning in medical image analysis is hampered by two challenges:

- heterogeneous raw data
- relatively small sample siz

For prostate cancer diagnosis, these two challenges can be conquered by using a tailored deep CNN architecture and performing an end-to-end training on 3D multiparametric MRI images with proper data preprocessing and data augmentation.

Attempts have been made to apply machine learning image analysis in clinical practice. Studies show that numerous use cases in clinical practice could be supported with machine learning. For example, on the basis of the Mura Dataset from the Stanford ML Group, it has been shown that baseline performance in detecting abnormalities on finger studies and equivalent wrist studies is on a par with the performance of radiologists. However, the baseline performance of convolutional networks comes in lower than that of the best radiologists in detecting abnormalities on the elbow, forearm, hand, humerus, and shoulder.

Numerous cases, including deepsense.ai's right whale recognition system, show that it is possible to tune a model enough to perform well on a limited dataset. Thus, the prospects for building models that outperform human doctors in detecting abnormalities are tantalizing.

An interesting practical example comes thanks to the paper a deep convolutional neural network-based automatic delineation strategy for multiple brain metastases stereotactic radiosurgery. Precise brain metastases targeting delineation is a key step for efficient stereotactic radiosurgery treatment planning. In the paper, an algorithm was used to segment brain metastases on contrast-enhanced magnetic resonance imaging datasets. Developing tools to support delineation of critical organs could save medical doctors a lot of time.

SUMMARY – FUTURE SAVINGS WITH AI

According to The Lancet, global healthcare spending is predicted to increase from $9.21 trillion in 2014 to $24.24 trillion in 2040. The spending is predicted to increase both in developing countries due to improving access to medical treatment, and in developed countries facing the challenge of providing care for their aging populations.

As a business, healthcare is unique because its provision is not measured solely by revenue. Potential savings and the ability to provide treatment for larger groups of people are better measures of the importance of AI to healthcare. According to Healthcare Global, AI is predicted to bring up to $52 billion in savings by 2021, enabling care providers to manage their resources better. A significant part will come from leveraging image recognition, as earlier diagnosis translates into lower treatment costs and greater patient well-being, as was clearly shown in this WHO study.

As modern radiology increases the adoption of machine learning to automatically interpret medical images and describe what they show, significant advantages will result, including including lower costs and further steps towards automating the diagnosis process.

STUDIES SHOW THAT NUMEROUS USE CASES IN CLINICAL PRACTICE COULD BE SUPPORTED WITH MACHINE LEARNING.

#1 TOP AI BLOG
TOP 50 BLOGS

aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning –
a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing
– Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

The State of Natural Language Processing – Giant Prospects, Great Challenges

By Pawel Godula,
Director of Customer Analytics,
deepsense.ai

Natural Language Processing (NLP) is one of the most dynamic areas of AI and provides business with a host of opportunities. Yet there are many challenges to overcome.

[Continued on Next Page]

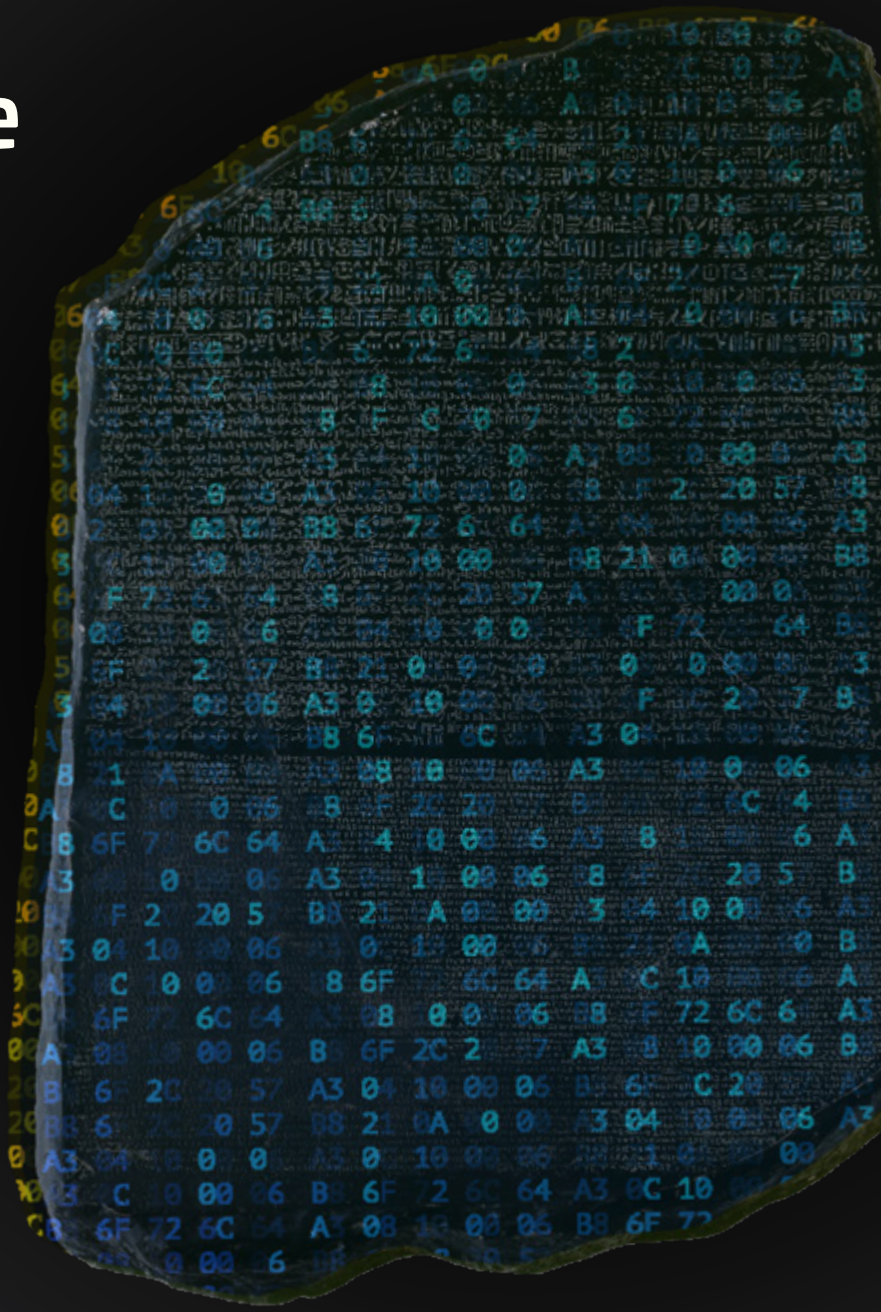

#1 TOP AI BLOG
TOP 50 BLOGS

aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning – a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing – Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

In 1950, Alan Turing proposed a test for a machine's intelligence. In order to pass it, a machine should be able to hold a conversation that would be indistinguishable from one a human would produce.

Putting aside machines' capacity to think, Turing considered the ability to respond appropriately and follow the conversation to be the point at which automation ends and intelligence begins. "If they find a parrot who could answer to everything, I would claim it to be intelligent being without hesitation," claimed Denis Diderot two hundred years before Turing came along.

The idea behind the Turing test implies that natural language processing is considered a frontier for Artificial Intelligence. History and the modern approach to text processing prove that the ability to analyze, understand and think are totally different.

HISTORY AND THE MODERN APPROACH TO TEXT PROCESSING PROVE THAT THE ABILITY TO ANALYZE, UNDERSTAND AND THINK ARE TOTALLY DIFFERENT.

WHERE NLP IS USED

On the most basic level, NLP doesn't require Artificial Intelligence and machine learning. The auto correction function found in most of today's text software is a simple, rules-based way to analyze natural language by comparing written text with a dictionary database. It improves countless lives every day, yet has nothing to do with AI.

Automatic filtering is another example of NLP, one used primarily by media vendors seeking to avoid the swearing and hate speech in the comments sections at the ends of articles. As machines can filter out only specific words, they are unable to understand the context or see through the simple tricks used to throw them off. Is "duck" a dirty word, after all?

NLP BUSINESS APPLICATIONS

Even with its flaws, today's NLP is maturing rapidly and getting increasing attention from business. Automated translation (AT) is the best benchmark of the current state of NLP. While far from perfect—it tends to confuse words and lose the context of the translation, often to comic effect—it is the best benchmark of the current state of NLP and is a very useful tool.

Naturally, thousands of people use NLP every day. There are mobile apps that can translate text from a video camera in real time. Just such as tool was put to use during a court hearing, when a British court failed to provide an interpreter for Mandarin speaker Xiu Ping Yang.

NLP is also used to enhance the customer experience, with chatbots being one of the most popular tools. According to a study done by Walker, 88% of buyers are willing to pay more for a better customer experience. That knowledge compelled Domino's Pizza to give its Dom chatbot a job helping customers order pizza through Facebook messenger or Google Home. Royal Bank of Scotland claims that the use of its mobile app increased 20% after a chatbot was incorporated.

But building a chatbot may not require that much Artificial Intelligence. Sometimes the chatbot provides only a predefined set of answers to the most common questions or automates processes like ordering a product or inquiring about opening hours.

REPUTATION MONITORING

Tasks like online reputation monitoring or sentiment analysis require not only the ability to identify particular words in a sentence, but to understand the sentence itself. deepsense.ai faced such a challenge while building a solution for the United Nations Office of Information and Communications Technology. The project was set up to uncover propaganda disseminated via Twitter. With fake news among today's most pressing concerns, both companies and governments need to constantly keep abreast of stories being published online. After all, fake stories outperform their non-fake counterparts in reaching audience on every subject, from business to science to technology. With NLP, it is possible to analyze not only the constant flow of language-based data from social media, but also to determine if the news shared is real or fake.

#1 TOP AI BLOG
TOP 50 BLOGS

aiworld CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning – a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing
– Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

THE LEGAL WAY

Another cutting edge example comes from the consultancy giants EY, PwC and Deloitte, which use NLP to review massive numbers of contracts or compliance with lease accounting standards. All that information — not only contracts and legal agreements but also emails, conversations and all text-based data — is unstructured. Gartner estimates that up to 80% of business data today is unstructured. NLP is one of the most powerful tools for analyzing it and gathering meaningful insights. As research done at the University of Rome Tor Vergata suggests, adopting NLP-based techniques results in 12% less effort spent on classifying equivalent requirements.

The NLP market has been forecast to reach $22.3 billion by 2025. The main driver behind the acceleration is building new, scalable solutions for real life problems, both for business and consumers.

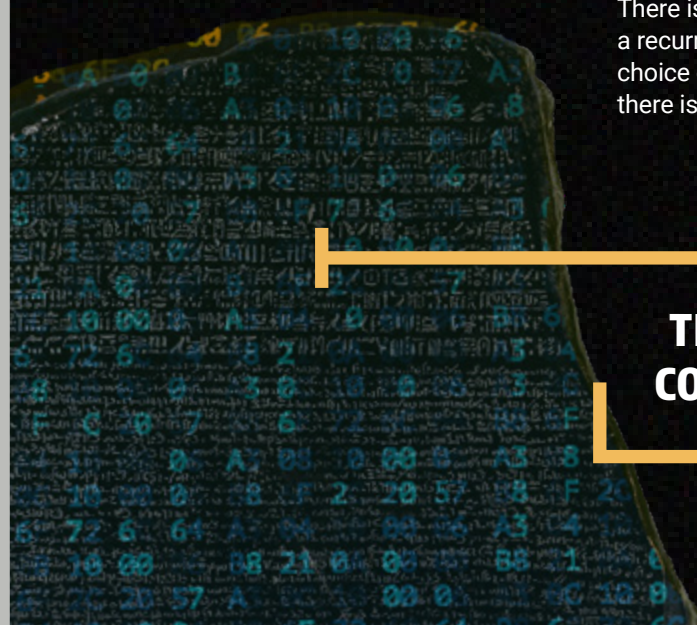

NLP – KEY CHALLENGES AND HOW IT WORKS

What sets NLP apart from other challenges is the nature of data. Unlike images that can be resized to 500×500 pixels, there is no way to standardize sentences to be "always seven words long". Natural language is processed by recurrent neural networks that pick words one by one in a queue, as sentences vary in length. Some algorithms work character-by-character while others process a group of characters at once.

Transfer learning uses pre-trained neural networks and is a huge challenge in NLP. So huge, in fact, that some consider it impossible. For images, on the other hand, the first layers are mostly responsible for finding the general structure of the image. Extracting information specific to the task tends to happen in the deeper layers. Recurrent nets, however, do not produce such a clear distinction – when using a pre-trained net, data scientists often "freeze" the first layers of the network during further training. There is no obvious way to freeze a part of a recurrent neural network. Furthermore, the choice of training data is vital in this task – there is a tangible difference between the language used in newspaper articles and, say, tweets, possibly an even bigger disparity than between photos and cartoon-style drawings. In NLP, context is a crucial factor.

The need to interpret everything in a given context is the greatest challenge NLP faces. Chatbots and automated assistants are gaining in popularity. According to PwC data, 72% of business execs use automated assistants and 27% of consumers weren't sure if their last customer service interaction was with a human or a chatbot. As chatbots are a convenient way to provide answers to the most common questions, they struggle to provide answers to questions that call for deeper context or to analyze more complex feedback. If the customer writes in social media that "this place is way too cool for me", are they referring to the temperature or to how groovy the place is?

These examples are only the tip of the language contextuality iceberg. Conversation commonly incorporates irony, puns and cultural references, which makes the message more human but blurry and unclear for machines. In NLP, the language has to be natural.

THE NEED TO INTERPRET EVERYTHING IN A GIVEN CONTEXT IS THE GREATEST CHALLENGE NLP FACES.

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning –
a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing
– Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

FUTURE PERSPECTIVES – WORD EMBEDDINGS AND GRAMMATICAL GENDER

The challenge of making words understandable to computers has been tackled by fasttext and word2ved models. In these models, every word on the internet is marked with a multidimensional vector marking its meaning and connotations (commonly known as word embedding). Vectors of a cat, a kitten and a kitty are similar and dissimilar from a car, computer or telephone. With that knowledge, it is easier to tune up the model.

The key difference between image recognition and NLP is that the former makes it possible to train a model using additional images, as our recent publication showed. With NLP, on the other hand, all the words in the network are already described, so there is no need to expand its knowledge. Instead, it is entirely up to the model and the data scientist's skill to perform tasks with the data provided. Considering that, word embeddings can be compared to the first layers of a pre-trained image recognition network.

Because of the highly contextualized data it must analyze, Natural Language Processing poses an enormous challenge. Language is an amalgam of culture, history and information, the ability to understand and use it is purely humane.

Other challenges are associated with the diversity of languages, with their morphology and flexion. Finnish grammar with sixteen noun cases is hard to compare with English. In many European languages, there is grammatical gender to deal with.

What's more, words' meanings tend to vary depending on not only the context of the rest of the text, but also the social background or lifestyle of a conversation partner. Is "savage rap concert" a positive or negative review? It depends on whether it was written by a teen or a grandma, so even proper identification of a language and understanding the whole sentence doesn't mean the interpretation will be good.

The Turing test proves that the ability to respond in a human way is not merely a question of understanding.

THE TURING TEST PROVES THAT THE ABILITY TO RESPOND IN A HUMAN WAY IS NOT MERELY A QUESTION OF UNDERSTANDING.

#1 TOP AI BLOG
TOP 50 BLOGS

aiworld
CONFERENCE & EXPO

SPONSOR'S LETTER

FEATURED ARTICLES

AI Customer Targeting Levels
From Customer Segmentation to Recommendation Systems

Reinforcement Learning –
a Moonshot or Today's Most Underhyped Technology

Machine Learning Advancing Medical Imaging and Analysis

The State of Natural Language Processing
– Giant Prospects, Great Challenges

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com