

# Strategic Technical Insights

---

## DEEP LEARNING FOR HUMAN DECISION SUPPORT

**Prepared for**

Alain Auger  
Corporate Office

**Prepared by**

Erica Wiseman,  
Intelligence Analyst,  
National Research Council of Canada / Government of Canada  
[Erica.Wiseman@nrc-cnrc.gc.ca](mailto:Erica.Wiseman@nrc-cnrc.gc.ca)

**DRDC Project #:** FE22071713

**NRC Project #:** EW16-04

**Report submitted:** January 20, 2017

March 2017

DRDC-RDDC-2017-C061

NRC-KM employees make every effort to obtain information from reliable sources. However, we assume no responsibility or liability for any decisions based upon the information presented.

The scientific or technical validity of this Contract Report is entirely the responsibility of the Contractor and the contents do not necessarily have the approval or endorsement of Department of National Defence of Canada.

© Her Majesty the Queen in Right of Canada, as represented by the Minister of National Defence, 2017

© Sa Majesté la Reine (en droit du Canada), telle que représentée par le ministre de la Défense nationale, 2017



[This page intentionally left blank]



National Research  
Council Canada

Conseil national  
de recherches Canada



Table of Contents

**1 Executive Summary .....5**

**2 Background .....7**

2.1 Context..... 7

2.2 Key Issues..... 7

2.3 Key Questions ..... 7

**3 Introduction .....8**

3.1 Defining Deep Learning..... 8

3.2 Temporal Distribution and Subjects ..... 10

**4 Deep Learning Architectures .....12**

4.1 Architectures in Full Dataset..... 16

4.2 Architectures used in Human Decision Support Subset ..... 17

**5 Emerging Trends.....18**

5.1 Emerging Trends: Full Dataset ..... 18

5.2 Emerging Trends: Human Decision Support Subset ..... 19

5.3 Quantum Computing ..... 32

**6 Key Players.....33**

6.1 Key Players: Full Dataset ..... 33

6.1.1 Canadian Areas of Expertise ..... 34

6.1.2 International Collaborations ..... 34

6.1.3 Canadian-International Collaborations..... 36

6.2 Key Players: Human Decision Support Subset ..... 36

6.2.1 Canadian Areas of Expertise ..... 38

6.2.2 International Collaborations ..... 41

6.2.3 Canadian-International Collaborations..... 43

**7 Application Areas .....44**

7.1 Future Applications..... 46

**8 Deep Learning Limitations .....48**

**9 Commercially-Available Technologies .....49**

**10 Conclusions .....50**

**11 References .....53**

**12 Appendix A: Attachments .....68**

**13 Appendix B: Methodology .....68**

13.1 Searches..... 68

13.2 Analysis ..... 69

13.3 R&D Momentum..... 69

## List of Figures

Figure 1. Publication Timeline.....	10
Figure 2. Top Subject Groups.....	11
Figure 3. Subjects in Human Decision Support Subset .....	12
Figure 4. Architectures, Full Set .....	16
Figure 5. Architectures by Task.....	17
Figure 6. Architectures Used in Subset Topics.....	18
Figure 7. Emerging Technologies in Deep Learning.....	19
Figure 8. R&D Momentum, Subset: Emerging & Brand New/Disappearing.....	20
Figure 9. Top Affiliations, Full Set .....	33
Figure 10. Top Canadian Affiliations, Full Set .....	34
Figure 11. Areas of Research, Canadian Affiliations, Full Set.....	34
Figure 12. International Collaborations, Full Set .....	35
Figure 13. Canadian-International Collaborations, Full Set .....	36
Figure 14. Top Affiliations, Human Decision Support Subset .....	37
Figure 15. Top Canadian Affiliations, Human Decision Support Subset .....	37
Figure 16. Canadian Affiliations Areas of Expertise .....	38
Figure 17. International Collaboration, Human Decision Support Subset.....	42
Figure 18. Canadian-International Collaborations, Subset .....	43
Figure 19. Applications in Subset.....	44
Figure 20. Applications in Subset by Human Decision Support Topics.....	45
Figure 21. IBM’s Projected Decision Support Market Opportunity.....	47

## List of Tables

Table 1. Summary of Findings.....	5
Table 2: Brief History of Deep Learning Research .....	8
Table 3. Category Definitions.....	11
Table 4. Deep Learning Architectures.....	13
Table 5. Commercial Deep Learning Technologies.....	49
Table 6. Attachments .....	68
Table 7. Search Terms.....	68

# 1 EXECUTIVE SUMMARY

In order to assist with long-term R&D planning and the prioritization of research topics, this scientometric study was commissioned to provide a high level overview of worldwide research activity in the field of deep learning for human decision support. This study will assist DRDC in uncovering and understanding the potential impact of new research on future defence and security capabilities and operations.

In all, 8,565 references to publications from 2011-2016 were retrieved and analyzed using text mining software and a variety of visualization tools to identify top architectures, emerging trends, key players, collaboration networks, application areas and limitations. The data was analyzed in two sets; the first included all records to provide an overview of research and trends in the broader deep learning field, while the second was a subset of the data, containing 2,395 publications, focused on human decision support. Additional web searching was executed to identify commercially available technologies. A summary of the findings is provided in Table 1.

**Table 1. Summary of Findings**

Key Question	Key Findings
<b>Definitions</b>	Deep learning is an application of machine learning (which itself is a branch of artificial intelligence) that is comprised of a set of architectures or algorithms based on (artificial) neural networks (and the neuroscience of how the human brain observes, analyses, learns and makes decisions). Deep learning algorithms can automate the extraction of complex representations (a.k.a. abstractions or features) without human intervention. They are based on a layered, non-linear, hierarchical network of learning and representing data, where higher-level (more abstract) features are defined in terms of lower-level (less abstract) features. Deep learning is a scientific development that provides more advanced performance than shallow learning architectures such as decision trees, support vector machines and case-based reasoning and is an important step towards artificial intelligence by making machines independent of human knowledge. <sup>1</sup> Deep learning has been applied to fields such as computer vision, automated speech recognition, natural language processing, medical diagnostics, social media, political forecasting, customer relationship management, recommendation systems and bioinformatics and is particularly relevant in an age of big data and the Internet of Things.
<b>Top Architectures</b>	<ul style="list-style-type: none"> <li>• Convolutional Neural Networks (primarily for imaging tasks)</li> <li>• Deep Neural Networks (primarily for speech related tasks)</li> </ul>
<b>Emerging Trends</b>	Human decision support subset’s emerging topics were: <ul style="list-style-type: none"> <li>• <i>Visualizations</i></li> <li>• <i>Surveillance</i></li> <li>• <i>Event detection</i></li> <li>• <i>Computer aided instruction</i></li> <li>• <i>Recommender systems</i></li> <li>• <i>Risk analysis</i></li> </ul> Very limited evidence was found in the dataset that suggested quantum computing would have a significant impact on deep learning, although research is currently exploring the potential.
<b>Key players</b>	Top international key players include: <ul style="list-style-type: none"> <li>• Chinese Academy of Sciences (445 publications/113 publications in subset)</li> <li>• Microsoft, International (284 publications/64 publications in subset)</li> <li>• Tsinghua University, China (207 publications/61 publications in subset)</li> </ul>

Key Question	Key Findings
	Top Canadian key players include: <ul style="list-style-type: none"> <li>• University of Toronto (85 publications/17 publications in subset- 2<sup>nd</sup> place)</li> <li>• Université de Montreal (67 publications/22 publications in subset –1<sup>st</sup> place)</li> <li>• University of British Columbia (34 Publications/12 publications in subset)</li> </ul>
<b>Application areas</b>	The top applications in which human decision support plays a role include: <ul style="list-style-type: none"> <li>• Diagnostics (medical) (263 publications)</li> <li>• Data mining (260 publications)</li> <li>• Searching (251 publications)</li> <li>• Big data analytics (214 publications)</li> </ul>
<b>Limitations</b>	In general, there are three main limitations to deep learning, including: <ul style="list-style-type: none"> <li>• Architectures and output models are, in many ways, seen as a black box,</li> <li>• Inability to reason or apply learning to another context,</li> <li>• Requires massive amounts of data.</li> </ul>
<b>Commercially available technologies</b>	Commercially available technologies were found for the following areas: <ul style="list-style-type: none"> <li>• Decision support</li> <li>• Security</li> <li>• Medical imaging/diagnostics</li> <li>• Surveillance</li> <li>• Emotion/sentiment recognition</li> <li>• Big data analytics</li> <li>• Event detection</li> <li>• Recommender systems</li> <li>• Risk analysis</li> <li>• Forecasting</li> <li>• Data fusion</li> </ul>

Recommendations for next steps include:

- Monitoring of quantum computing technologies to remain abreast of potential developments of interest.
- Explore and track issues and developments in deep unsupervised learning.
- Explore the connection between IBM’s Cognitive Solutions and DRDC’s interest in human decision support.
- Exploration of reinforcement learning, which takes a slightly different tack than most of the algorithms discussed in this report.

## 2 BACKGROUND

### 2.1 Context

In order to assist with long-term R&D planning and the prioritization of research topics, scientometric studies are being commissioned to provide a high level overview of worldwide research activity in scientific domains. These studies will assist DRDC in uncovering and understanding the potential impact of new research on future defence and security capabilities and operations.

Deep learning, also known as deep structured learning, hierarchical learning or deep machine learning, is an application of machine learning (which itself is a branch of artificial intelligence) that is comprised of a set of architectures or algorithms based on (artificial) neural networks (and the neuroscience of how the human brain observes, analyses, learns and makes decisions). Deep learning algorithms can automate the extraction of complex representations (a.k.a. abstractions or features) without human intervention. They are based on a layered, non-linear, hierarchical network of learning and representing data, where higher-level (more abstract) features are defined in terms of lower-level (less abstract) features. Deep learning is a scientific development that provides more advanced performance than shallow learning architectures such as decision trees, support vector machines and case-based reasoning and is an important step towards artificial intelligence by making machines independent of human knowledge.<sup>[1]</sup> Deep learning has been applied to fields such as computer vision, automated speech recognition, natural language processing, medical diagnostics, social media, political forecasting, customer relationship management, recommendation systems and bioinformatics and is particularly relevant in an age of big data and the Internet of Things.

This scientometric study will focus on deep learning within the context of human decision making and support. This study will not look at research on applications of deep learning for adaptive behavior of machines, robots, autonomous systems or cars.

### 2.2 Key Issues

The objective of this study is to detect and categorize the international R&D domains in the field of deep learning for human decision support, as well as to identify the commercially available products on the market. Results of this project will be used to identify domains that could present an area of interest for which DRDC may wish to develop expertise and to support the selection of future deep dive projects.

### 2.3 Key Questions

1. Deep Learning
  - a. What is deep learning?
  - b. What are the different scientific approaches to deep learning within the context of human decision making?
2. Emerging trends
  - a. What are the emerging trends in deep learning research? (selected topics will be described in detail)
  - b. Is there any evidence that quantum computing might change deep learning in the future?

---

<sup>[1]</sup> Najafabadi MM, Villanustre F, Khoshgoftaar TM, Seliya N, Wald R, Muharemagic E. Deep learning applications and challenges in big data analytics. *Journal of Big Data*. 2015;2(1):1-21.

3. Key players
  - a. Who are the key players (organizational level)?
  - b. What are their areas of expertise?
  - c. What are their collaboration networks (both internationally and in Canada)?
4. Application areas
  - a. What are the current application areas for deep learning within the context of human decision making?
  - b. What future applications for deep learning for decision making might be possible based on the literature and market reports (e.g. real-time decision making; reduced data-to-decision time)?
5. What are the documented or projected limitations of deep learning within the context of human decision making?
6. What are the commercially-available technologies that use deep learning to support human decision making?

### 3 INTRODUCTION

#### 3.1 Defining Deep Learning

The field of deep learning, which in essence is a branch of artificial intelligence that tries to emulate how the human brain automatically extracts features and abstractions from data to learn,<sup>1</sup> has its roots in work on neural networks. The concept of neural networks has been theorized since the 1940’s but much of the early research was highly theoretical and could not be put into practice yet. Recent research, along with dramatically faster graphical processing units (GPUs) and lowered cost of computing hardware, has led to significant advances in deep learning’s ability to solve AI problems that have challenged the community for years.<sup>2,3,4</sup> Table 2 presents a brief history of the key developments that have advanced deep learning research to where it is today.

**Table 2: Brief History of Deep Learning Research**

Year	Major Event
1943	Warren McCulloch (University of Illinois) and Walter Pitts (University of Chicago) create a computational model for neural networks based on mathematics and algorithms called threshold logic.
1958	Frank Rosenblatt (Cornell Aeronautical Laboratory) creates the perceptron, an algorithm for pattern recognition based on a two-layer computer neural network using simple addition and subtraction. He also proposed additional layers with mathematical notations, but these wouldn’t be realized until 1975.
1980	Kunihiko Fukushima (NHK Broadcasting Science Research Laboratories, JP) proposes the Neoconitron, a hierarchical, multilayered artificial neural network that has been used for handwriting recognition and other pattern recognition problems.
1989	Scientists were able to create algorithms that used deep neural networks, but training times for the systems were measured in days, making them impractical for real-world use.
1992	Juyang Weng and colleagues (University of Illinois) publishes Cresceptron, a method for performing 3-D object recognition automatically from cluttered scenes.

<b>2004</b>	Geoffrey Hinton gets a small amount of funding from the Canadian Institute For Advanced Research (CIFAR) to form the Neural Computation and Adaptive Perception program – an invite only group of researchers working on creating computer systems that mimic organic intelligence.
<b>2006</b>	The term “deep learning” begins to gain popularity after a paper by Geoffrey Hinton and Ruslan Salakhutdinov (University of Toronto) showed how a many-layered neural network could be pre-trained one layer at a time.
<b>2009</b>	NIPS Workshop on Deep Learning for Speech Recognition discovers that with a large enough data set, the neural networks don’t need pre-training, and the error rates drop significantly.
<b>2012</b>	Artificial pattern-recognition algorithms achieve human-level performance on certain tasks. Google’s deep learning algorithm discovers cats.
<b>2014</b>	Google buys UK artificial intelligence startup Deepmind for £400m.
<b>2015</b>	Facebook puts deep learning technology - called DeepFace - into operations to automatically tag and identify Facebook users in photographs. Algorithms perform superior face recognition tasks using deep networks that take into account 120 million parameters.
<b>2016</b>	Google DeepMind’s algorithm AlphaGo masters the art of the complex board game Go and beats the professional Go player Lee Sedol at a highly publicized tournament in Seoul.

Sources: Marr, 2016;<sup>5</sup> Hernandez, 2014.<sup>6</sup>

The study of deep learning sits at the intersection of research in the fields of neural networks, graphical modelling, optimization, pattern recognition and signal processing.<sup>4</sup> Deep learning has been considered as either a class of machine learning techniques<sup>4</sup> or representation-learning methods<sup>2</sup> based on a set of algorithms that use many hierarchical layers of non-linear processing units for pattern classification, feature extraction and transformation. Deep learning algorithms are comprised of deep architectures of consecutive layers. The objective of deep learning algorithms is to stack up the non-linear transformation layers in order to learn complicated abstract representations of the data in a hierarchical manner by passing the data through multiple layers. Each layer applies a non-linear transformation on raw inputs and provides representations (a.k.a. features or abstractions) at a higher, more abstract level as an output.<sup>1,2</sup> Each successive layer uses the output from the previous layer as input. The final output is usually a prediction of some properties you are trying to discern from the inputs, for example, whether an image is a picture of a cat or not. These algorithms are considered as “deep” based on the number of parameterized transformations (i.e., a processing unit that has trainable parameters such as weights) a signal encounters as it is processed through multiple hidden layers, between the input and the output layer.<sup>3</sup> The more layers the signal passes through, the more complicated the non-linear transformations are, resulting in a final representation that is a highly non-linear function of the input data.<sup>1</sup> Deep learning algorithms may be supervised or unsupervised.<sup>3</sup>

In laymen’s terms, for the task of face recognition, a deep learning algorithm learns the edges in different orientations in the first layer; in the second layer it composes the edges to learn more complex features like different parts of a face such as lips, noses and eyes, regardless of small variations in the positions. In the third layer it uses these features to learn even more complex features such as the shape of different faces. These final representations can be used in face recognition applications. However, not all deep learning algorithms are used to reconstruct a pre-defined sequence of representations at each layer (e.g., edges, eyes, faces), as final outputs might be used to build classifiers or for data indexing using relational or semantic knowledge.<sup>1,2</sup>

Two key properties of modern deep learning algorithms that make them so successful are 1) the generative nature of the architectural model which typically requires adding an additional layer on top

to perform classification tasks, and 2) an unsupervised pre-training step that makes use of large amounts of unlabeled training data to automatically extract structures and regularities in the input features and reduce the common risk of overfitting in small datasets.<sup>4,2</sup> By adding multiple (e.g., 5-20) non-linear layers to the algorithm, it can increase both the selectivity and the invariance of the representations. Basically, the algorithm can thus be both selective of minute details, allowing it to distinguish a Samoyed dog from a white wolf, and at the same time be insensitive to large, irrelevant variations such as background, pose, lighting and surrounding objects.<sup>2</sup> Deep learning architectures can also “generalize in non-local and global ways, generating learning patterns and relationships beyond immediate neighbours in the data”.<sup>1</sup> Furthermore, unlike with other machine learning approaches which require feature engineering, deep learning eliminates this challenging manual step by automatically extracting complex representations and identifying the useful features for learning.<sup>1,3</sup> Finally, deep learning algorithms are functional with a wide variety of data types, from audio, seismic, image to natural language and is particularly well suited to large, messy, real-world, big data sets.<sup>3</sup>

### 3.2 Temporal Distribution and Subjects

To address the key questions in this project, a search on deep learning and its associated architectures was conducted in *Scopus*, *Inspec* and *NTIS*. A list of the specific search terms used is provided in Table 7. The initial search resulted in a dataset of 8,565 publications from 2011-2016. Figure 1 presents the publication timeline and reveals that there has been a significant jump in publications in the last two years.

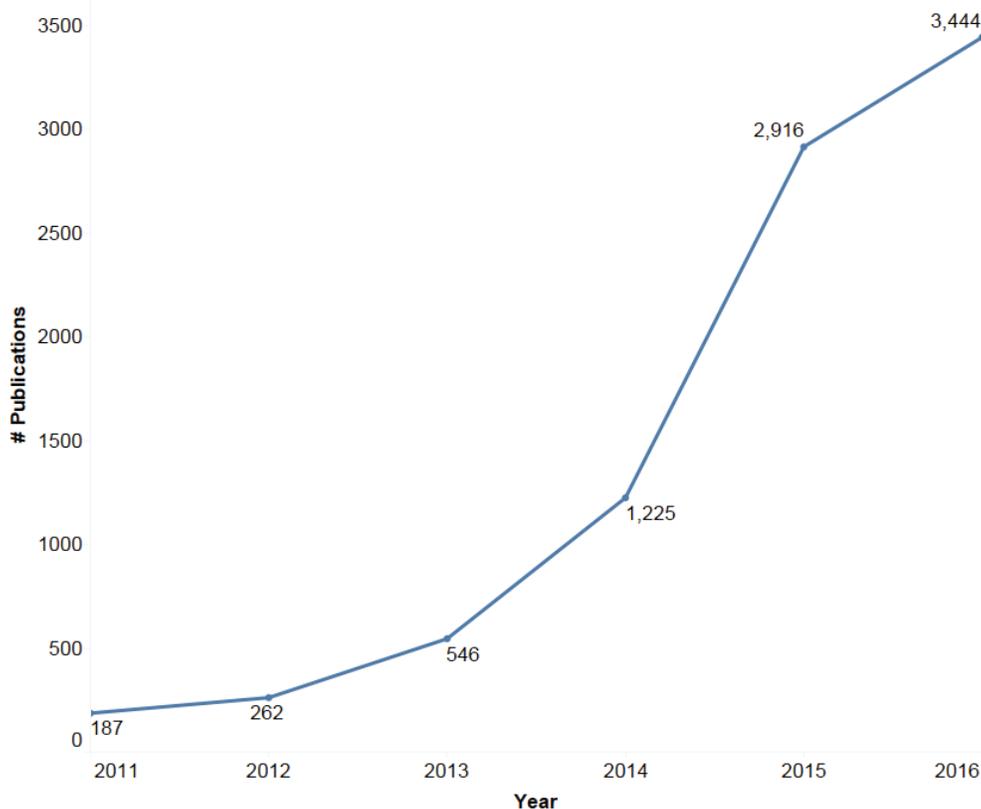


Figure 1. Publication Timeline

To identify the major research areas within the field, the terms in the various keyword fields and those derived from natural language processing of the title were merged and assigned to 99 subject groups covering 8,042 publications (94% of the full dataset). Figure 2 presents the top subjects, colour coded by their respective categories. The categories are defined in Table 3:

Table 3. Category Definitions

Category	Definition
Application	The general field in which the deep learning architectures are being used (e.g. <i>Health</i> )
Architecture	The scientific approach to deep learning (e.g. the broader class of algorithms, such as <i>Convolutional neural network</i> )
Data type/source	Data type or source that is being computed by the deep learning architecture (e.g. <i>Audio data</i> )
General	General concepts to the field of artificial intelligence (e.g. <i>Learning systems</i> )
Task	The activity the deep learning architecture is being used to support (e.g. <i>Image analysis &amp; processing</i> )

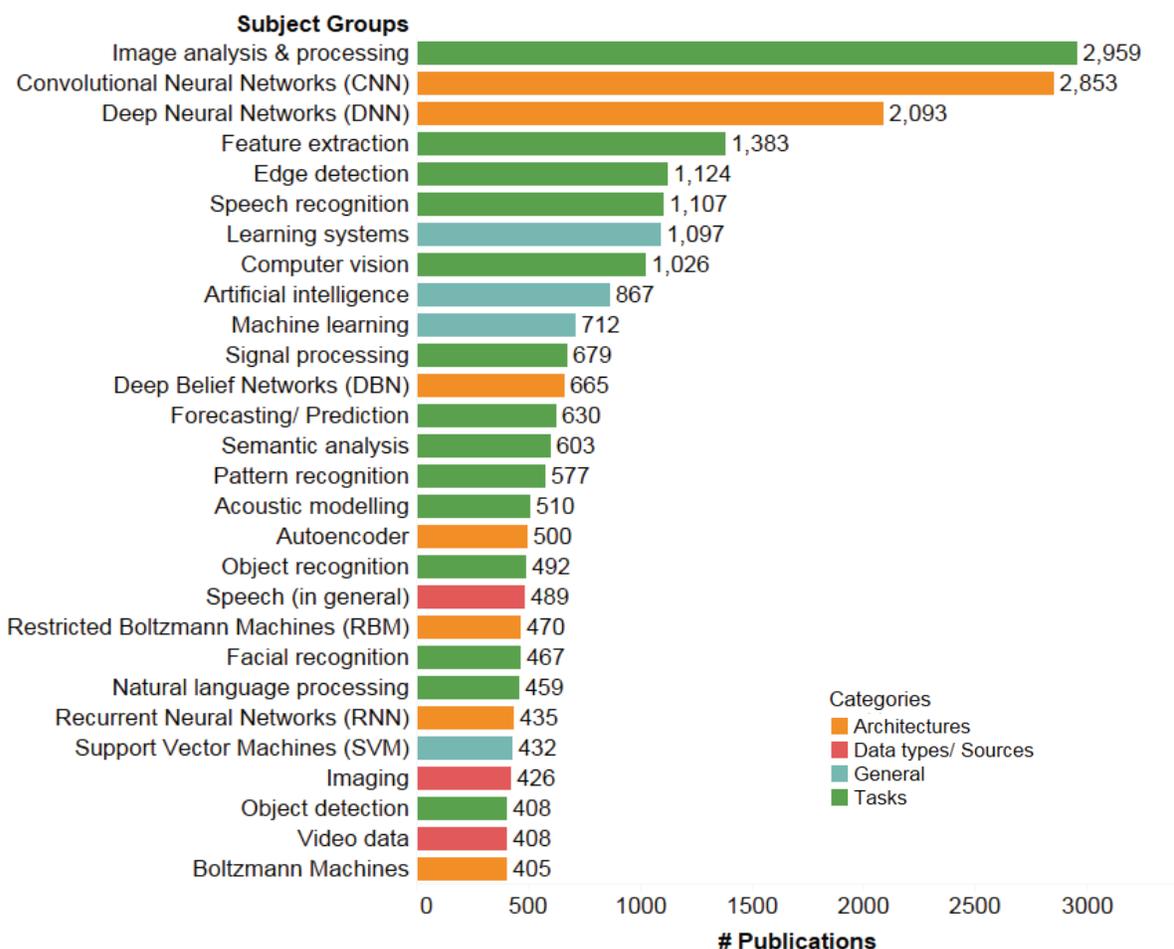


Figure 2. Top Subject Groups

This report provides analysis both of the broader deep learning field, as well as more specifically on a subset of the field that is focused on human decision support. This subset, containing 2,395 publications from 2011-2016, was selected by the client following a review of the subject groups in the broader field. The selected subject groups in the human decision support subset are presented in Figure 3.

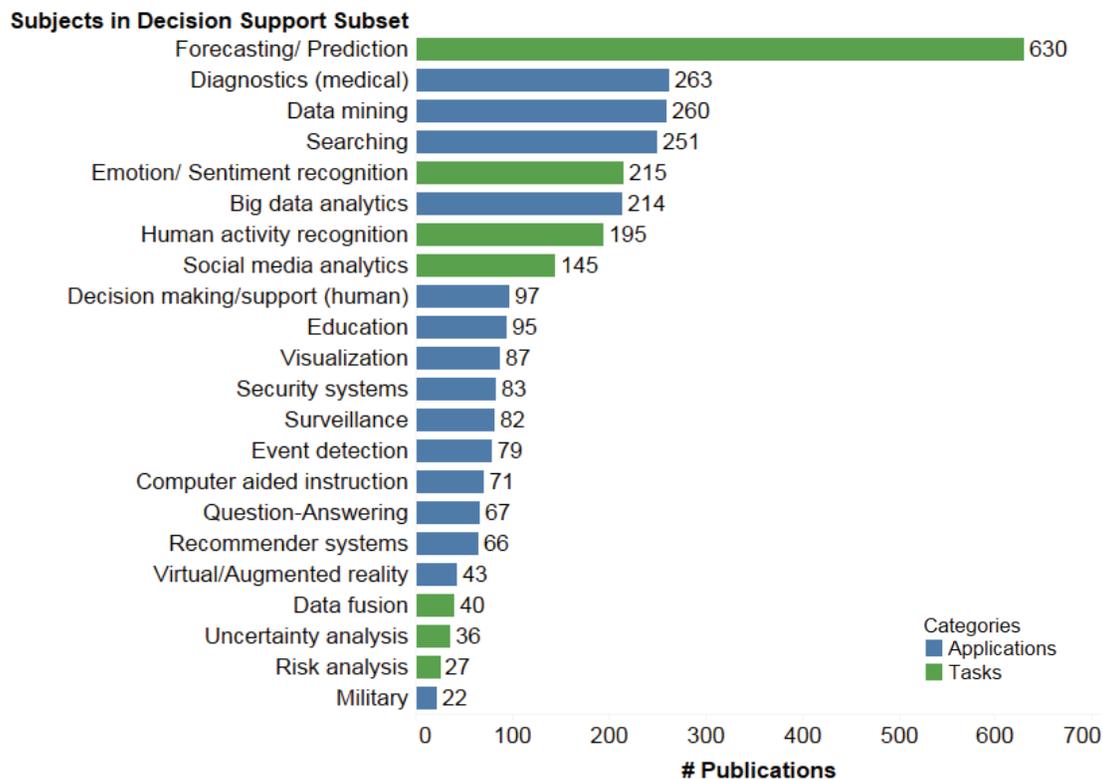
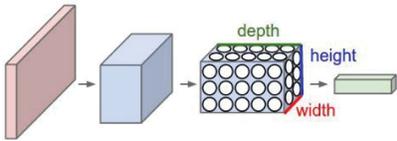
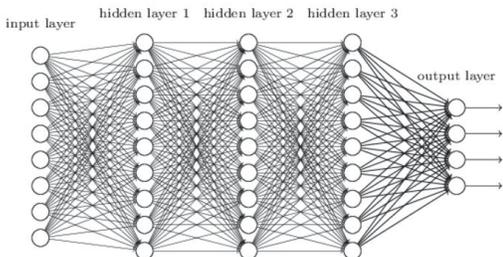
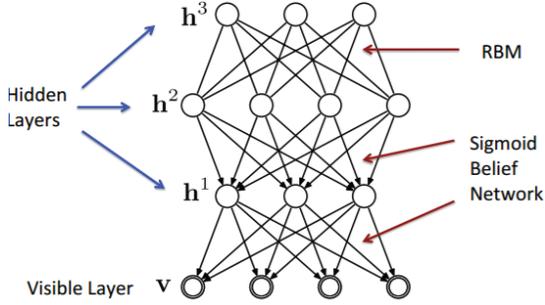
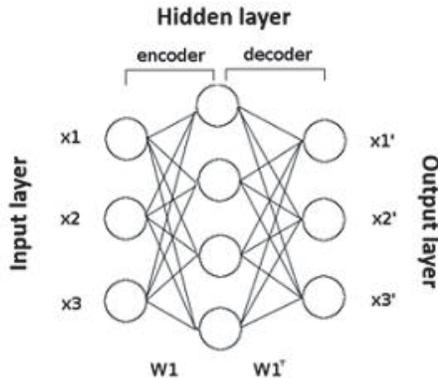


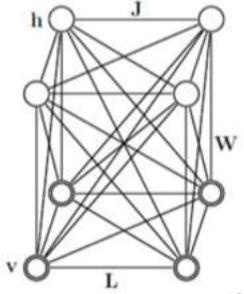
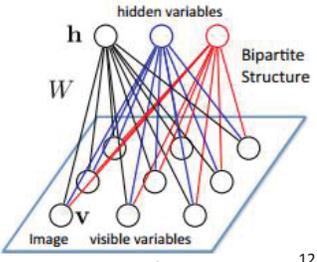
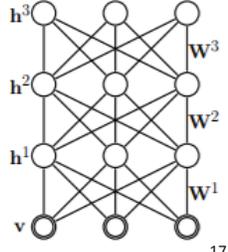
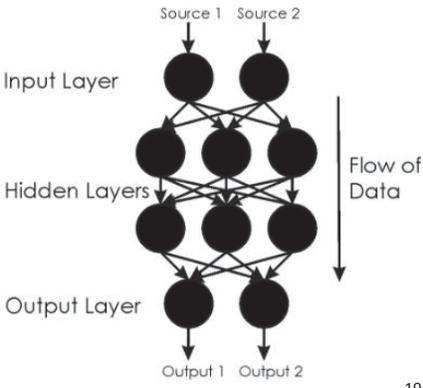
Figure 3. Subjects in Human Decision Support Subset

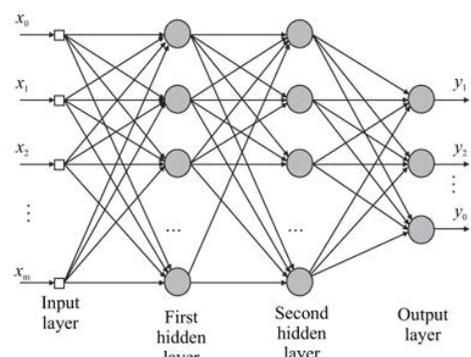
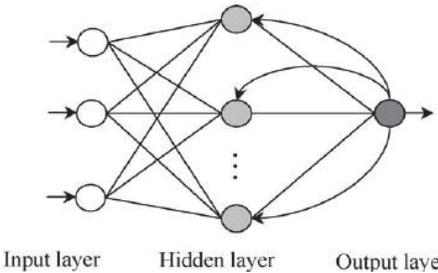
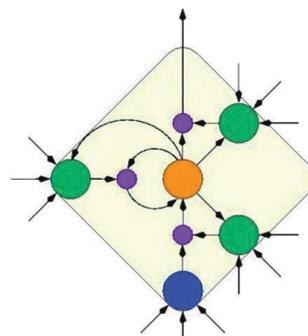
## 4 DEEP LEARNING ARCHITECTURES

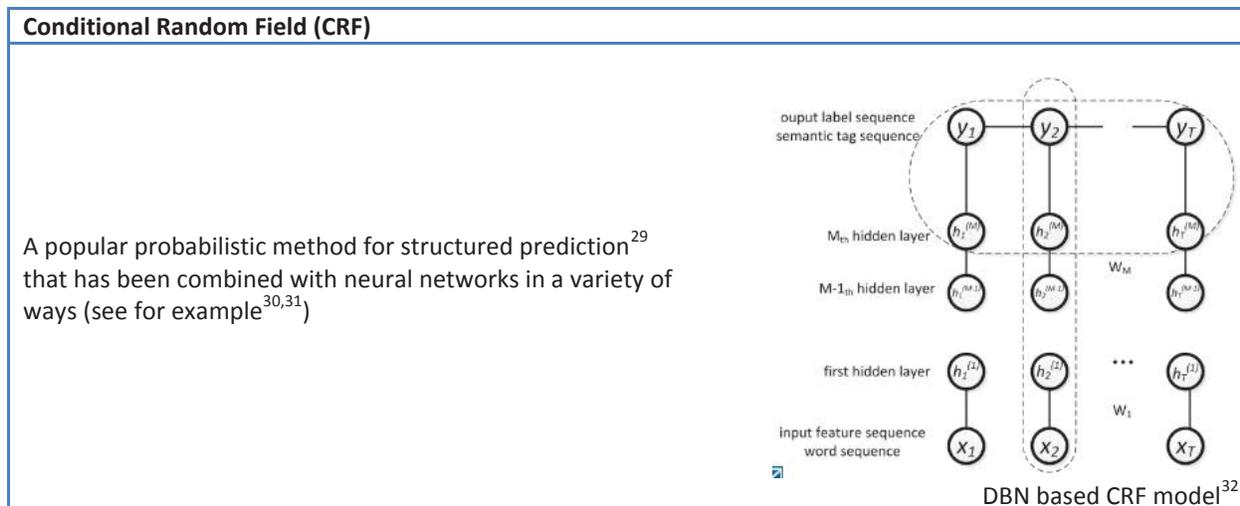
Deep learning architectures are numerous and often a branch or variant of some original parent architecture and include a multitude of different associated algorithms. With the field growing at such an incredible pace recently, new architectures, variants and algorithms appear every few weeks. It is important to note that comparing the performance of multiple architectures is challenging as their capabilities are often reported based on different datasets.<sup>3</sup> Table 4, which is an expanded version of a similar table presented by Deng 2014,<sup>4</sup> provides a brief overview and description of some of the major deep learning architectures, the year in which the architecture was first described (with reference to the original paper when possible), and a figure depicting the architecture.

Table 4. Deep Learning Architectures

Architecture/Description	Introduction Date/Figure
<p><b>Convolutional Neural Networks (CNN)</b></p> <p>A CNN is composed of one or more convolutional layers with fully connected layers (matching those in typical artificial neural networks) on top. It also uses tied weights and pooling layers.<sup>3</sup> It arranges its neurons in three dimensions (width, height, depth) and transforms 3D inputs into a 3D output.<sup>8</sup></p>	<p>1970s; seminal paper 1998<sup>7</sup></p>  <p>Image source<sup>8</sup></p>
<p><b>Deep Neural Network (DNN)</b></p> <p>A multilayer network with many hidden layers, whose weights are fully connected and are often initialized (pre-trained) using stacked Restricted Boltzmann Machines (RBM) or Deep Belief Networks (DBN). (In the literature, DBN is sometimes used to mean DNN).</p>	<p>Mid -2000s<sup>9</sup></p>  <p>Image source<sup>10</sup></p>
<p><b>Deep Belief Network (DBN)</b></p> <p>A stacked RBM. Probabilistic generative models composed of multiple layers of stochastic, hidden variables. The top two layers have undirected, symmetric connections between them. The lower layers receive top-down, directed connections from the layer above.</p>	<p>2006<sup>11</sup></p>  <p>Image source<sup>12</sup></p>
<p><b>Deep Autoencoder</b></p> <p>A DNN whose output target is the data input itself, often pre-trained with DBN or using distorted training data to regularize the learning.</p>	<p>1986<sup>13</sup></p>  <p>Image source<sup>14</sup></p>

<p><b>Boltzmann Machine (BM)</b></p> <p>A network of symmetrically connected, neuron-like units that make stochastic decisions about whether to be on or off.</p>	<p>1986<sup>15</sup></p>  <p>Image source<sup>16</sup></p>
<p><b>Restricted Boltzmann Machine (RBM)</b></p> <p>A special BM consisting of a layer of visible units and a layer of hidden units with no visible-visible or hidden-hidden connections.</p>	<p>1986<sup>13</sup> originally known as <i>Harmonium</i>. Popularized as RBM by Hinton in the mid 2000s</p>  <p>Image source<sup>12</sup></p>
<p><b>Deep Boltzmann Machine (DBM)</b></p> <p>A special BM where the hidden units are organized in a deep layered manner, only adjacent layers are connected, and there are no visible-visible or hidden-hidden connections within the same layer.</p>	<p>2009<sup>17</sup></p>  <p>Image source<sup>17</sup></p>
<p><b>Feedforward Neural Networks</b></p> <p>A network of perceptrons arranged in layers, with the first layer taking in inputs and the last layer producing outputs. The middle layers are hidden. Each perceptron in one layer is connected to every perceptron on the next.<sup>18</sup></p>	 <p>Image source<sup>19</sup></p>

<p><b>Multilayer Perceptron (MLP)</b></p> <p>A type of feed-forward neural network model that has one layer or more of hidden units and non-linear activations.<sup>22</sup> MLP networks are typically used in supervised learning problems and solved with backpropagation algorithms or less frequently for unsupervised learning using auto-associative structures.<sup>23</sup></p>	<p style="text-align: right;">1974,<sup>20</sup> 1986<sup>21</sup></p>  <p style="text-align: right;">Image source<sup>24</sup></p>
<p><b>Recurrent Neural Networks (RNN)</b></p> <p>A network containing at least one feed-back connection or loop allowing it to do temporal processing and learn sequences. The simplest form of a fully RNN is an MLP with the previous set of hidden unit activations feeding back into the network along with the inputs. RNNs have memory that allow them to learn sequences.<sup>25</sup></p>	<p style="text-align: right;">1991</p>  <p style="text-align: right;">Image source<sup>26</sup></p>
<p><b>Long Short-Term Memory (LSTM)</b></p> <p>An efficient, gradient based variation of RNN that overcomes challenges of storing information over extended time intervals (long-time-lag tasks) by enforcing constant error backflow.<sup>27</sup></p>	<p style="text-align: right;">1997<sup>27</sup></p>  <p style="text-align: right;">Image source<sup>28</sup></p>



### 4.1 Architectures in Full Dataset

Figure 4 presents the top architectures found in the full data set. Convolutional Neural Networks and Deep Neural Networks far outnumber all other architectures.

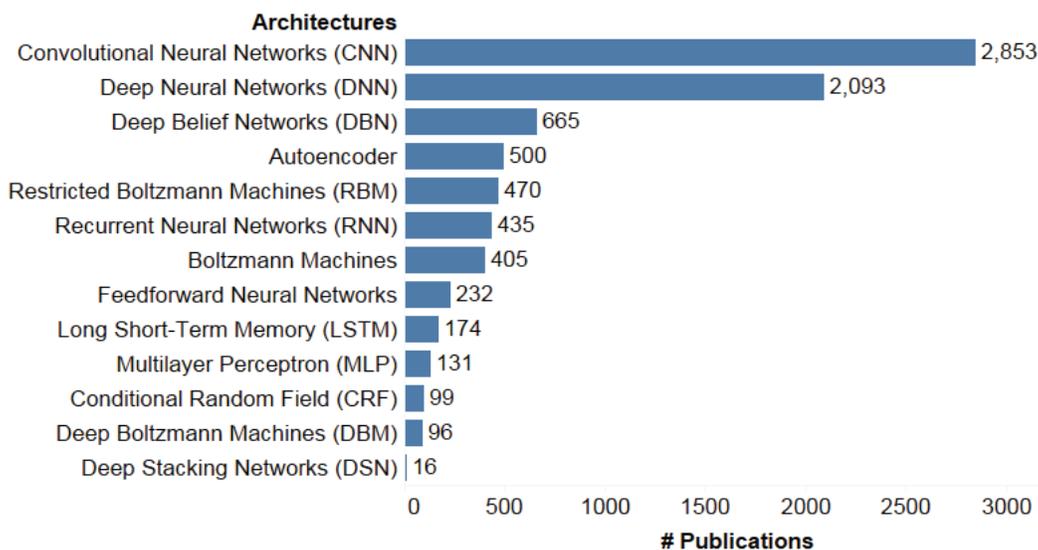


Figure 4. Architectures, Full Set

Figure 5 reveals the most frequent tasks for which each architecture is used. In this figure alone, *Imaging* includes *Image analysis & processing*, *Computer vision*, *Edge detection*, and *Feature extraction*; *NLP (speech)* includes *speech recognition*, *speech (as a data source)* and *acoustic modelling*; *NLP (general)* includes *Natural language processing (in general)* and *Semantics*; *General* includes *Learning systems*, *Artificial intelligence*, *Unsupervised learning*, and *Signal processing*. CNN is predominately used for imaging related tasks, whereas DNN, LSTM and arguably RNN are focused on speech related tasks.

Architectures	Subjects (group)			
	Imaging	NLP (speech)	NLP (general)	General
Convolutional Neural Networks (CNN)	3,542			
Deep Neural Networks (DNN)	443	1,466		
Autoencoder	392			281
Deep Belief Networks (DBN)	328	98		104
Feedforward Neural Networks	257			
Restricted Boltzmann Machines (RBM)	217			158
Boltzmann Machines	213			114
Recurrent Neural Networks (RNN)	100	161	72	
Conditional Random Field (CRF)	77	17	29	
Multilayer Perceptron (MLP)	67	29		22
Deep Boltzmann Machines (DBM)	58			33
Long Short-Term Memory (LSTM)	36	133		
Deep Stacking Networks (DSN)	17			3

Figure 5. Architectures by Task

## 4.2 Architectures used in Human Decision Support Subset

In addition to the broader themes identified in Figure 5, Figure 6 presents a closer look at the architectures used in each of the subject groups in the human decision support subset. Here, we see that CNN is the most frequently used architecture by all subject groups except in *Recommender systems* and *Uncertainty analysis* which uses DNN with a slightly greater frequency. Also worth noting is that CNN and DNN are nearly equally used in *Forecasting/Prediction, Data mining, Searching, Education, and Decision making/support (human)* and so these topics should be monitored to determine if there really is a preferred architecture.

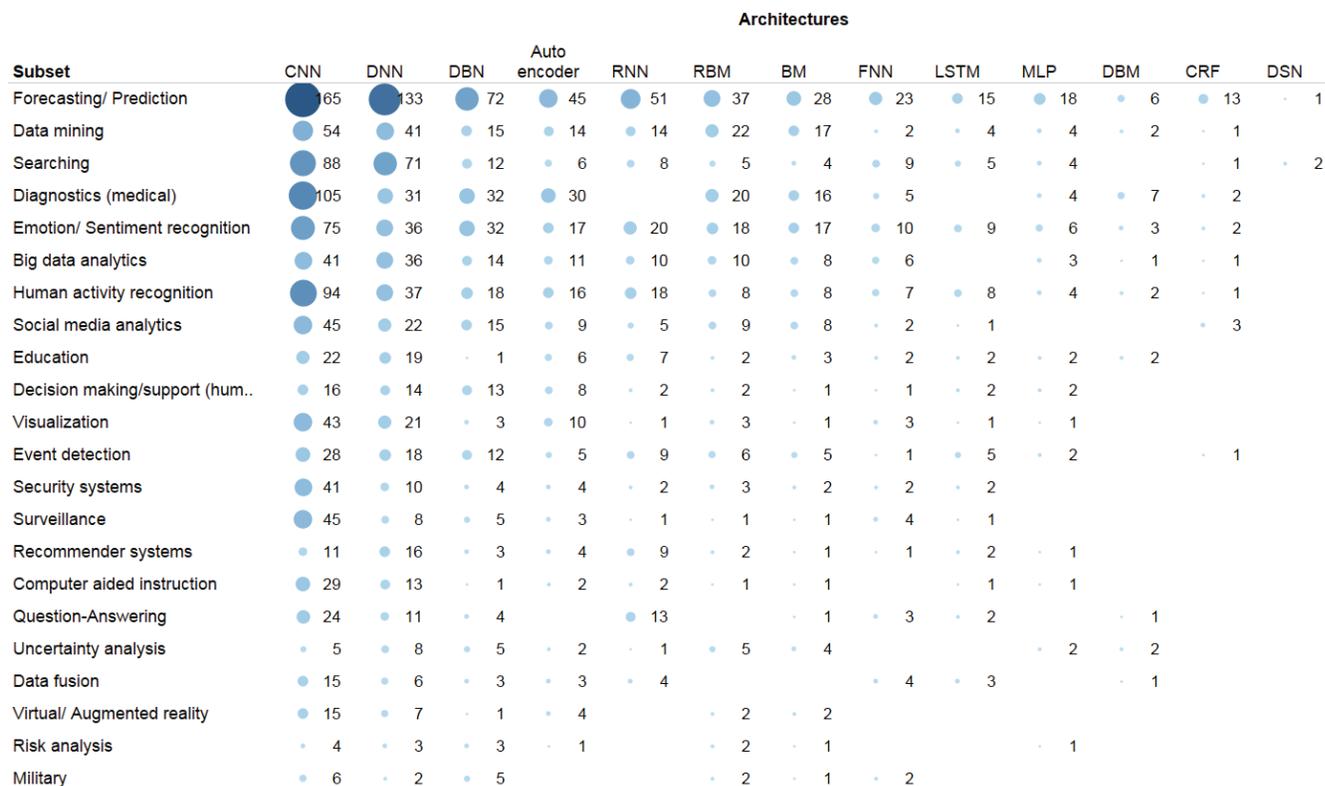


Figure 6. Architectures Used in Subset Topics

## 5 EMERGING TRENDS

Although a count of publications within each subject group provides some insight into the quantity of research activity on a given topic, it does not sufficiently convey the momentum of a topic relative to others in the same set of data. In other words, the quantity of records in a group does not necessarily suggest which topics are "hot", which are fading and which are emerging. For topics in the literature, the determination of momentum is based on the relative velocity of each subject group in the dataset. Further explanation of the methodology behind the momentum indicator is included in Section 13.3, but essentially it plots the standard deviation of standardized measures of publication counts and velocity (the rate of publication increase) on two axes. Nodes which plot to the left of the Y-axis intersection have below-average velocity, and those found below the X-axis have relatively smaller publication counts. A third dimension is added by sizing nodes relative to the total number of underlying publications. Even a small node which plots to the lower/right side of the axes may be of interest, since emerging topics are typically small in numbers as they begin to attract research attention and increase in velocity. The full R&D Momentum chart for both the full dataset and the human decision support subset are available in the Tableau workbook provided with this report.

### 5.1 Emerging Trends: Full Dataset

Figure 7 presents the emerging research subjects in the full dataset. The majority (44) of the subject groups in the full dataset fall into the Emerging quadrant, with 25 in both the Brand new/disappearing

and Hot quadrants and only five in the Established quadrant. This is not surprising given the significant jump in publications in the past two years alone as most research in this field should still be considered emerging.

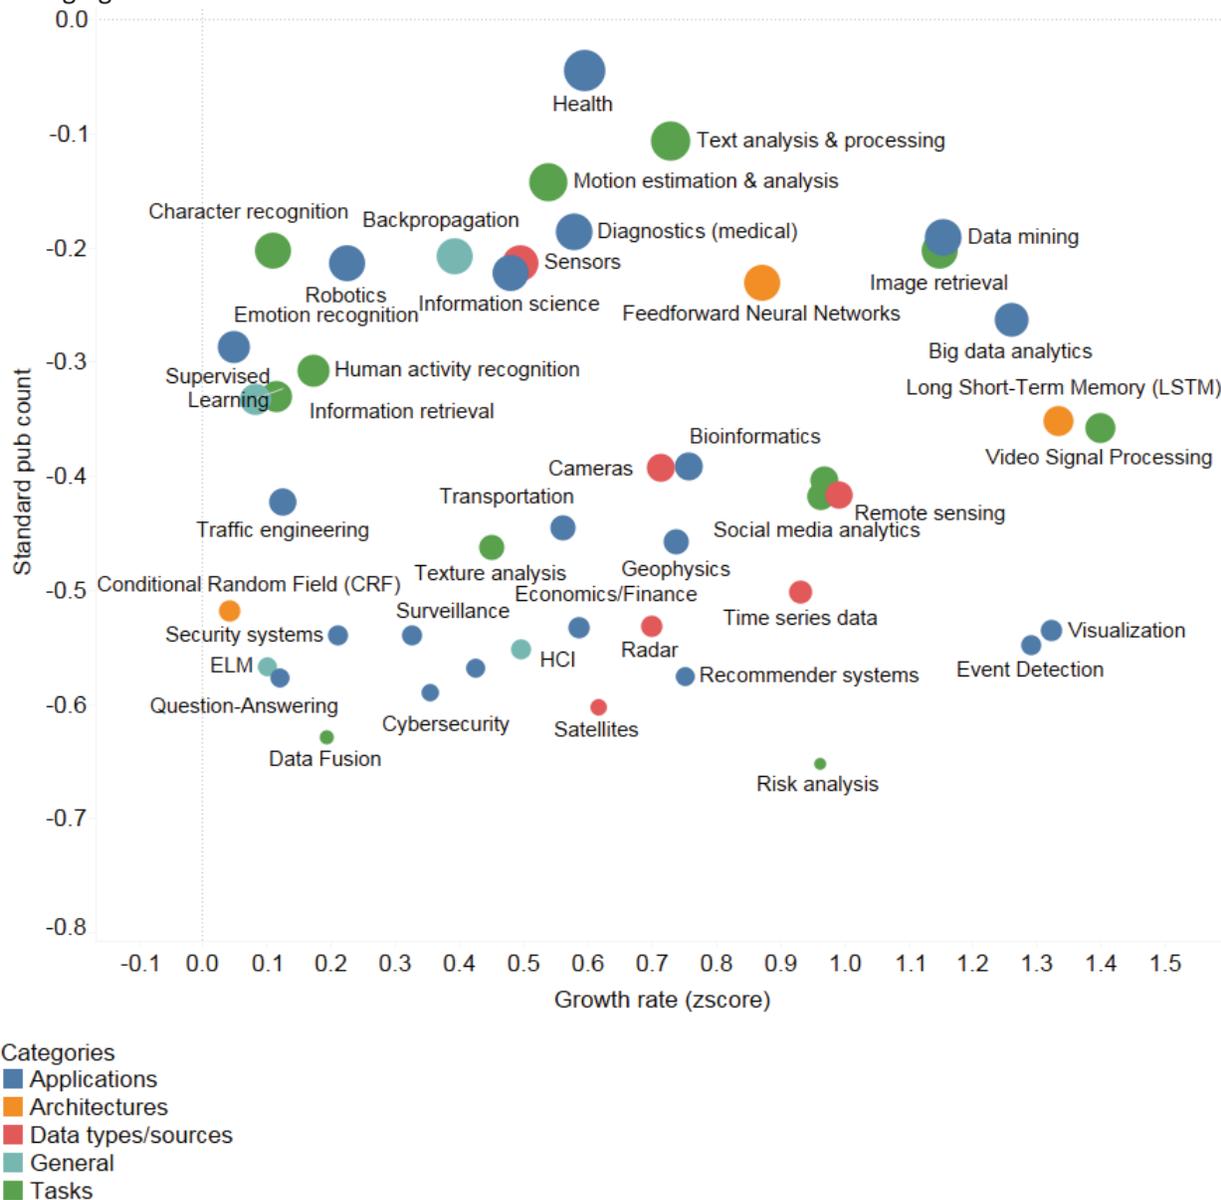


Figure 7. Emerging Technologies in Deep Learning

## 5.2 Emerging Trends: Human Decision Support Subset

Figure 8 presents the emerging and brand new/disappearing subject groups in the human decision support subset. Because the R&D momentum indicator is a measure of relative velocity, Figure 8 is calculated exclusively on the subset subject groups. As such there is some movement in the positioning of nodes compared to Figure 7, for example, *Data fusion* and *Question-Answering* are now in the Brand new/disappearing quadrant as opposed to the Emerging quadrant as seen above. A detailed description of the research in the Emerging quadrant and selected topics in the Brand new/disappearing quadrant is presented below.

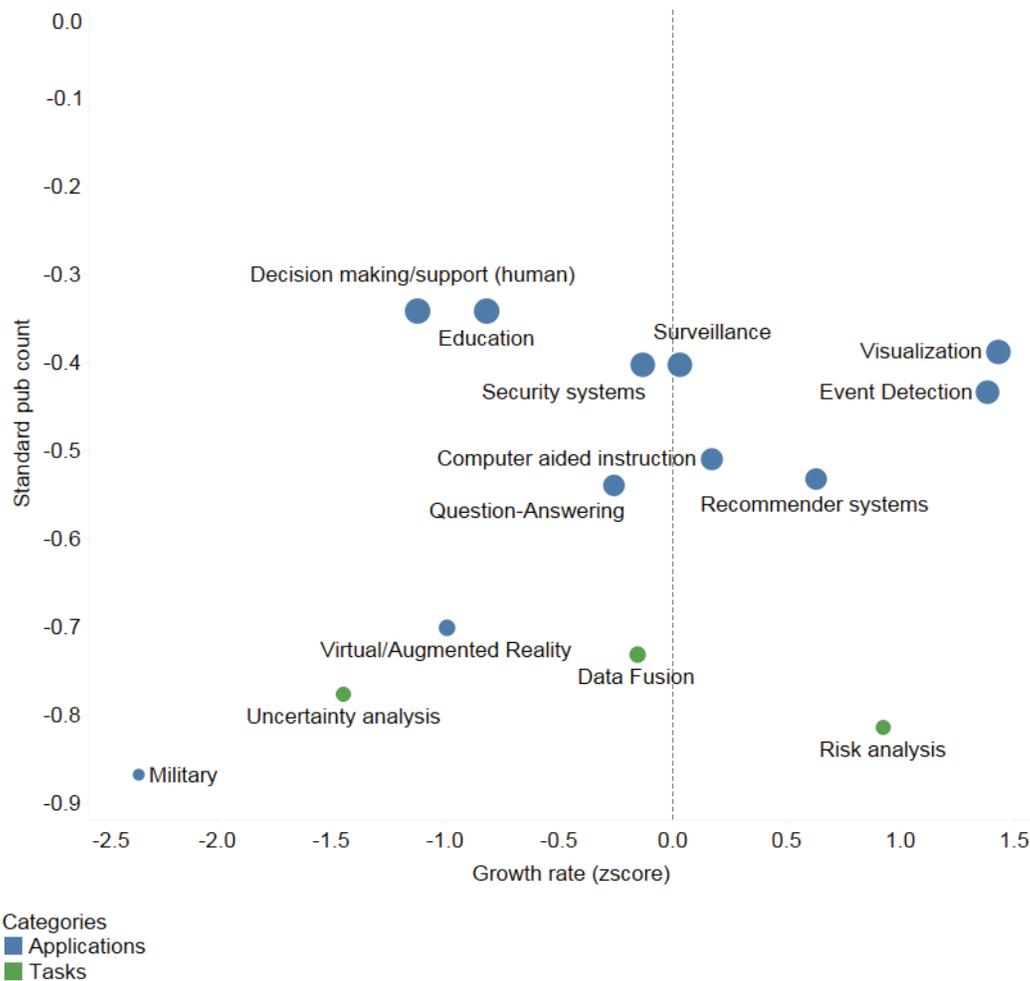


Figure 8. R&D Momentum, Subset: Emerging & Brand New/Disappearing

**Visualization**

The *Visualization* subject group has 87 publications, with 84% in 2015-2016 and 53% in 2016 alone. The largest group (22 articles) is focused on using visualization techniques to open the “black box” of deep learning, followed by 17 articles that discuss flow graphing (a.k.a. flow visualization) in relation to deep learning. Other smaller, but interesting, topics include seven articles on feature visualization methods and two on intelligence analysis. Some of the other, more generic topics in the group include visual classification and other computer vision related tasks.

While deep learning methods have proven to be effective at surpassing current state of the art machine learning algorithms in a wide variety of tasks, how they actually work is still somewhat unknown. In particular, how the deep learning algorithms actually learn, what precisely is happening in the hidden layers, or how the algorithm arrives at a particular classification or recognition decision, remains unclear.<sup>33-35</sup> One attempt to open the black box is to visualize the networks and the neural excitation in the layers.<sup>34</sup> This task is sometimes known as “deep visualization”<sup>36-38</sup> and includes a variety of approaches. For example, the Deep Genomic Dashboard (Deep GDashboard or DeMo Dashboard), which is used for transcription of factor binding and DNA sequencing, uses saliency mapping, temporal analysis

and class-specific visualization to visualize CNN, RNN and CNN-RNN to understand how the architectures make predictions.<sup>39</sup> The Sensitivity-characterised Activity Neorogram (SCAN) is another approach for understanding the inner workings of a DNN by analysing and visualising the sensitivity patterns of the neuron activities. SCAN produces a low dimensional visualization space to show neuron activity in a meaningful and interpretable way. This visualization space can be used to compare neurons, both within the same DNN and across DNNs that are trained for the same tasks.<sup>40</sup> Other approaches include a visualization technique based on deconvolutional networks to understand how deep learning algorithms choose features to classify objects.<sup>41,42</sup> Some have suggested that the use of a layer-wise relevance propagation algorithm provides a better qualitative and quantitative explanation of how algorithms arrive at their classification decisions than deconvolution methods, or sensitivity-based approaches, which typically produce visualizations in terms of a heat map.<sup>33</sup> Others have proposed a modified code inversion algorithm called feature map inversion to provide a deeper understanding of the inner mechanisms of what a filter extracts from images when deep learning algorithms are used.<sup>37</sup> Still others have suggested using prediction gradients as opposed to visualizing units and classes in pixel space to understand how neural networks encode concepts into individual units.<sup>43</sup> An extendable, easy to use, open source library for visualizing CNNs, called FeatureVis library for MatConvNet, is a general taxonomy and library of some of the main approaches to visualize the features learned by CNN networks.<sup>35</sup> One noted limitation of visualizing CNNs is that, in general, the various techniques do not provide accurate information about the stability of the algorithms.<sup>44</sup>

Data flow visualization, or data flow graphing, is a graphical representation of the flow of data through a system, showing what data is used as input and what is received as the output.<sup>45</sup> TensorFlow, a popular open source software library developed by Google's Machine Intelligence Research Group for deep neural networks research, uses data flow graphs. Three examples of flow visualization techniques that are used to improve our understanding of deep learning algorithms include "inversion, in which the aim is to reconstruct an image from its representation, activation maximization, in which we search for patterns that maximally stimulate a representation component, and caricaturization, in which the visual patterns that a representation detects in an image are exaggerated".<sup>46</sup> In this project's dataset, flow visualization has been used to investigate the dimensionality reduction ability of auto-encoder to determine if, when stacked, it can contribute to the success of deep learning tasks.<sup>47</sup> Flow visualization has also been used to reveal how a new CNN architecture with fewer parameters can more effectively classify traffic signs than the winning and runner-up CNNs in the German Traffic Sign Benchmark Competition. The visualizations showed that the new CNN learned pictographs of the signs and ignored the shape and colour information.<sup>48</sup> Another article used flow visualization to identify the actual parameters needed in deep learning algorithms to process an image with human vision-like quality.<sup>49</sup> Finally, a last article presented a deep learning programming framework based on heterogeneous architecture which can be used to build deep learning models through a visual interface and helps overcome several challenges related to designing deep learning architectures.<sup>50</sup>

### Surveillance

The *Surveillance* subject group has 82 publications, close to 70% of which are from 2015-2016 with 50% from 2016 alone. Based on the R&D momentum diagram in Figure 8 the application of deep learning to surveillance is quite new and is only just starting to grow at a notable rate. The bulk of the scientific research has tested one deep learning architecture or another (although often CNN) to perform the task in question, whether it is detection, recognition, identification, or classification, with better than current state-of-the-art (i.e. non-deep learning algorithm) performance results.

The majority of the 2015-2016 articles are focused on object/event detection and human detection in surveillance videos. To date, the majority of video surveillance footage in the U.S. is not being monitored by a human and is simply a vast network of evidence collection which is only reviewed after an event. Deep learning technologies, which facilitate the processing and analysis of vast amounts of video footage, often in real-time, to detect not only humans but also changes in poses and behavior, is a technology area that has recently seen significant investment and research.<sup>51,52</sup>

Research on object/event detection covers such topics as:

- the detection of bags (purses, backpacks),<sup>53</sup>
- the classification of objects<sup>54</sup> or image sets<sup>55</sup> from video surveillance,
- distinguishing fine-grained object categories in low resolution images<sup>56</sup>,
- taking advantage of temporal and spatial features in videos to improve event and action detection<sup>57</sup> even in real-time scenarios,<sup>58</sup> and
- efficient indexing and retrieving of objects of interest from large-scale surveillance videos.<sup>59</sup>

Other studies have looked at human action recognition,<sup>60</sup> with one in particular investigating the use of deep learning and multiscale symbolic time series analysis in dynamic data-drive application systems (DDDAS) for real time activity recognition, with applications to border surveillance scenarios, and situation awareness in Intelligence Surveillance and Reconnaissance (ISR) missions.<sup>61</sup> Only one study on rapid event detection concluded that no single CNN worked best for all events, attributing some of the challenge to the differences between object-driven events vs scene-based events.<sup>62</sup>

Human detection is discussed in terms of pedestrian detection from aerial platforms<sup>63</sup> or traffic surveillance,<sup>64</sup> and studies on human detection in surveillance have looked at identifying visual attributes or “soft-biometrics” (e.g. gender, clothing)<sup>65-67</sup> and the challenges of night time detection.<sup>68</sup> Human detection is also concerned with the re-identification problem, which involves detecting a single individual across multiple camera views.<sup>69-71</sup> While deep learning approaches to the problem have achieved remarkable performance, one recent study by Liu et al., highlighted the need to go beyond current approaches of extracting discriminative features from the whole frame at one glimpse and move toward examining multiple, highly discriminative regions of the image in detail through multiple glimpses.<sup>69</sup> Concerns about de-identification, i.e. concealing a person’s identity, have also been raised.<sup>72,73</sup>

Other research is focused on scream/sound detection,<sup>74</sup> health surveillance,<sup>74,75</sup> challenges in facial recognition both in terms of blur, dramatic pose variation and occlusions,<sup>76</sup> as well as challenges with night-time surveillance,<sup>77</sup> vehicle detection and annotation,<sup>78</sup> license plate detection,<sup>79</sup> vehicle re-identification,<sup>80</sup> and finally, crowd surveillance<sup>81</sup> and in particular crowd counting.<sup>82</sup>

Two examples of companies that are using deep learning in surveillance tools include Camio and Herta Security. [Camio](#), a U.S. intelligent video monitoring start-up company, has integrated deep learning technology to learn user preferences in advance to selectively upload relevant footage to their cloud. This technology is also able to identify the most important events in the data and send personalized relevant alerts based on past experience. The algorithm allows for differentiation between ordinary and anomalous events as well as the analysis of data based on keywords.<sup>83</sup> [Herta Security](#), a Spanish biometric security software company, uses deep learning in their facial recognition systems that deliver some of the world’s fastest results, particularly in crowded environments and can successfully deal with occlusions, changes in facial expression and difficult lighting conditions. Their software work both in real-time and “faster than” real time for forensic purposes.<sup>84</sup>

### Event detection

The *Event detection* subject group has 79 publications, 85% in the last two years, and includes research on intrusion detection (17 publications), anomaly detection (17) as well as more general event detection (42) including multimedia and acoustic events. Although the *Surveillance* subject group touched on event detection these two groups are fairly mutually exclusive with only 8 publications overlapping between the two groups. Furthermore, the surveillance articles focused more on real-time video analysis, whereas the event detection articles are more focused on video clips than streaming video. Of the 64 articles between 2015-2016, 40 report on the use of various deep learning algorithm for various kinds of event detection and recounting, and report performance improvements beyond current state-of-the-art 73% of the time (with positive results the remainder of the time). Interestingly, some of the existing state-of-the-art that is being surpassed by deep learning algorithms is other, previously tested, deep learning approaches. The performance improvements are reported in terms of precision rates (mean average precision or MAP scores),<sup>85,86</sup> accuracy,<sup>87,88</sup> and processing speed (compared to humans).<sup>89</sup> In addition, some articles also report reductions in training time or improvements in training performance using new approaches.<sup>85,90</sup>

The bulk of the studies in this group are focused on either acoustic event detection (16) or network intrusion detection (12). Acoustic event detection (AED) is sometimes a part of the broader activity of multimedia event detection which seeks to combine multiple features (e.g. appearance, colour, texture, motion and audio) from multi-modalities<sup>91</sup> and has applications in automated surveillance, machine hearing and auditory scene understanding.<sup>92</sup> The AED studies aim to recognize non-speech acoustic events (e.g. door knock, crowd cheering, engine noise) in real-world noisy scenarios,<sup>93</sup> by classifying audio frames among a set of semantic units, sometimes called “noisemes”.<sup>94</sup> Other approaches have been to use video-level labels to train audio classification tasks.<sup>90</sup> In general, deep learning approaches are proving effective for AED, even in situations with high levels of surrounding noise<sup>95</sup> and are more effective when detail rich features (such as spectrograms) are extracted.<sup>96,97</sup>

In terms of network intrusion detection systems (IDS), while many machine learning methods have been applied, there is a continued need to improve detection performance and accuracy.<sup>98</sup> This goal is increasingly being met through deep learning approaches. For example, one recent study used a traffic-based IDS built on a novel randomized data partitioned learning model (RDPLM) that achieved 99.984% accuracy and 21.38s training time on a benchmark dataset with results that outperformed other well-known machine-learning models such as sequential minimal optimization, deep neural networks, reduced error pruning tree and random tree.<sup>87</sup> In another similar study, a deep belief network model, using restricted Boltzmann machine and a supervised back-propagation algorithm, was found to outperform support vector machine and neural network approaches to intrusion classification.<sup>99</sup>

Other research in this group focused on detecting anomalies in crowd behavior,<sup>100,81</sup> energy consumption,<sup>101</sup> medical based signals (e.g. ECGs or other physiological data),<sup>102</sup> climatic events,<sup>103</sup> or disaster response.<sup>104</sup> Two articles of potential interest from a military perspective were found. The first discussed the use of deep learning (CNN) to perform pattern recognition tasks in digital 2D Ground Penetrating Radar (GPR) images of buried explosive hazards thus reducing the burden on highly trained field operatives.<sup>105</sup> The second article described the use of CNN feature extraction, as opposed to traditional bag-of-visual-words, to classify the semantics of Tweet images to detect social events in real time from Twitter.

One of the challenges that are discussed in both the *Event detection* and *Surveillance* subject groups is the need to extract features from both the spatial and temporal dimensions of video data for general surveillance or event detection. In particular, the lack of spatial and temporal cues leads to challenges with large variations in lighting, viewing angles, body poses and occlusions.<sup>69</sup> Deep learning techniques are providing promising results,<sup>106</sup> and in some studies, even outperforming previous state-of-the-art approaches.<sup>69,81,107,108</sup>

Many of the studies in the *Event detection* group are concerned with testing the deep learning approaches on pre-existing datasets. Nine studies on video based event detection used the [TrecVid](#) database while 11 of the network intrusion studies used the [KDD99](#) dataset or the more recent and improved [NSL-KDD](#) dataset.

### Computer aided instruction

The *Computer aided instruction* group has 69 articles, 53 of which are published in 2015-2016 and 77% in 2016 alone. Of the 41 articles published in 2016, 23 relate to medical imaging, 16 to disease detection and five which are specific to cancer. This overlap with medical imaging is largely because the term *Computer aided instruction* is an engineering controlled vocabulary term that has been used to index some (but not all) research that is using deep learning to improve medical image analysis, likely the body of research that is more closely related to computer assisted diagnosis (CAD). Other topics covered in this group include language learning (2) and student assessment (1). Because of the dominance of imaging related research, CNNs were the most commonly discussed deep learning architecture (24 articles) whereas DNN was only discussed in four articles.

Deep learning architectures have been found to be more effective and accurate, in medical imaging, at a variety of tasks, including, but not limited to:

- Cell segmentation<sup>109</sup>
- Lung cancer diagnosis through CT slices<sup>110</sup>
- Classifying benign versus malignant breast masses<sup>111</sup>
- Detection and quantification of blood brain barrier permeability in dynamic contrast-enhanced MRIs<sup>112</sup>
- Extracting, from segmented lesion masks, latent features that can predict short-term disease activity in patients with early MS symptoms more accurately than lesion volume (the common imaging biomarker)<sup>113</sup>

Three of the medical imaging articles make specific reference to including some sort of decision mechanism<sup>114</sup> along with the CNN in terms of understanding a diagnostic decision of a CAD<sup>115</sup> or simply to support clinical decision making.<sup>116</sup>

In terms of language learning, deep learning architectures have been used in pronunciation verification systems,<sup>117</sup> such as the Computer Aided Language Learner (CALL) system.<sup>118</sup> A deep-learning based model has also been used for automated writing assessment as part of an evidence-centered design framework that provides accurate prediction of student science learning.<sup>119</sup>

### Recommender systems

The *Recommender systems* group has 66 publications, 52 of which were published in 2015 and 2016. The 2016 reviewed articles (33) showed that, when specified, 36% of recommendation systems were

based on textual information such as research papers, including in the medical field,<sup>120</sup> movie plot summaries, news articles or blog posts<sup>121</sup> followed by 21% on visual data (images or videos),<sup>122</sup> as well as smaller amounts on website click-throughs<sup>123</sup> and music.<sup>124</sup>

Recommendation systems (RS) have used a variety of techniques to reach their end goal, including one of the most popular approaches: collaborative filtering (CF). However CF has a number of limitations including data sparsity (i.e. when users rate only a limited number of items) and cold start problems (i.e. the difficulty in bootstrapping the RSs for new users or new items).<sup>125-127</sup> One of the issues with the data sparsity is the problem of feature transformation or feature learning. On the one hand, projection methods can find new representations of data, but not in non-linear data or very sparse datasets. On the other hand, unsupervised feature learning with deep neural networks is generally fragile to non-Gaussian noises, outliers and high dimensional sparse data.<sup>128</sup> Some have tried to use auxiliary information, through document modelling, to boost CF, however this suffers from the inherent limitations of the bag-of-words model and leads to poor understanding of contextual information.<sup>127,129</sup> User response prediction models have also been used as an RS technique, in particular for web-based personalized RSs, but these are limited to either linear models (which loses the ability to explore feature interactions) or require manually building up high-order combination features (which results in a heavy computation in the large feature space).<sup>130</sup> Previous work has used topic models or averages of word embedding to map text to factors to support RS based on latent factor models in order to overcome the sparsity and cold-start problems.<sup>121</sup>

All of these approaches, however, have been surpassed by the addition of the deep learning architectures, resulting in better accuracy performance. For example, RNNs have been used to encode text sequences into latent vectors to recommend scientific papers,<sup>121</sup> the *Ask Me Any Rating* content-based recommender system uses DNN for information retrieval,<sup>131</sup> a hybrid RS for Netflix based on deep learning neural networks and CF provides good prediction for cold start items and performs better than four existing RS models for non-cold start items,<sup>126</sup> Deep Bidirectional Long Short-Term Memory showed significant improvement over state-of-the-art methods for music recommendations,<sup>124</sup> DNNs, leveraged with three feature transformation methods (factorisation machines (FMs), restricted Boltzmann machines (RBMs) and denoising auto-encoders (DAEs)) have surpassed state of the art models with web-based personalized RS,<sup>123</sup> and RNNs, in combination with feedforward networks, have shown significant improvements in real-time e-commerce recommendation systems.<sup>132</sup>

Further improvements are found with context aware models which enable sensing and analysis of user context to provide more personalized recommendations to users.<sup>133</sup> Often, context-aware recommender systems have increased sparsity and therefore tend to only incorporate a small set of pre-defined explicit contexts that don't necessarily accurately reflect the user context. However, the use of unsupervised deep learning techniques, to learn environmental features such as? low dimensional unsupervised latent contexts, has been showing improvements in recommendation accuracy.<sup>134</sup> Kim et al. describe a novel context-aware recommendation model called [ConvMF](#) that uses CNN and probabilistic matrix factorization to capture contextual information in documents to enhance rating prediction accuracy. This model not only significantly outperforms the state-of-the-art recommendation models, even when rating data is extremely sparse, but also successfully captures subtle contextual differences of words in documents.<sup>129</sup> Autoencoders have also been used to learn patterns of item-related contextual behavior to predict the category of items that should be recommended to users in a given context.<sup>133</sup> Citation and quote recommendation systems can also be improved with a combination of a variety of algorithms, including CNN, which has been used to capture important local semantic

features of text, and RNN, which reflect the ordering of sentences and words in the context of text, and which outperformed the best state-of-the-art systems by 46.7%.<sup>135</sup>

Recommender systems have shown promise for military and intelligence applications as they have the potential to shift computational support from being reactive to being predictive. In particular, in terms of cybersecurity, RS can be used to generate prioritized lists for defence actions, to detect insider threats, to monitor network security and to expedite other analyses.<sup>136</sup> The Australian Department of Defence is developing recommender systems for command and control of autonomous vehicles as part of a project for developing intelligent multi-UxV planners with adaptive collaborative/control technologies. The goal of the RS component is to reduce the cognitive burden on operators through recommendations, alerts and constraints on top of lower-level autonomy. In particular, they are aiming to develop RS that can *learn* recommendations at a range of command and control levels and that work when heuristics are not known or in new (previously unprepared for) contexts. Part of the *learning* comes from DeepMind's (Google) Deep Q Reinforcement Learning algorithm, the same supervised deep learning and reinforcement learning technique that was used to teach a computer to play Atari games.<sup>137</sup>

There are unique challenges for the military in using RS that are not faced by commercial counterparts. While both are dealing with high dimensional, large volumes of data, "ground truths" are not ultimately known by the defence community and successful use of an RS is harder to quantify (i.e. can't be measured by increase in sales). The decisions that are being made in a military context are based on situations in which past activities do not necessarily reflect future ones (adversaries and threats are continually evolving) and the consequences of decisions can be significantly more serious than in the commercial world. As such, the need for transparency in recommendations is also higher and it may be necessary to incorporate a confidence measure with each associated recommendation. Future challenges for military and intelligence usage of RS include establishing user trust, preserving privacy and security, adapting to user environment, developing multilevel metrics, and promoting system extensibility.<sup>136</sup>

### Data fusion

The *Data fusion* group has 40 publications, with 33 published in 2015 and 2016. These records discuss both unimodal and multimodal data fusion from a variety of data types including image and video. Data fusion is increasingly of interest as new devices, with decreasing data storage costs and increasing data collection capacity, flood the market. However, the analytic techniques needed to make sense of this big data are not fully developed. Deep learning techniques have been showing some promise in this field. But combining deep learning and data fusion for such tasks as event prediction, error reduction and data compression, still carries with it many of the classic questions about data fusion including what is the most effective way to combine data from different modalities and does the fusion method affect the performance with different classifiers.<sup>138,139</sup> Tao et al recently explained that data fusion is usually performed prior to classification to reduce dimensions and the complexity of the classification model to improve classification performance.<sup>140</sup>

Many of the articles in the past two years present various deep learning and data fusion frameworks within in a variety of scenarios. For example, using simulated and real-world environmental data, a stacked auto-encoder extracted key features from multimodal sensor data to compute compact representations that were used in analytic tasks.<sup>139</sup> In another study using real world events, temporal multimodal data from audio, visual, depth and sensors were combined for classification tasks using a

hybrid recurrent neural network that could learn both modality specific temporal dynamics and dynamics in a multimodal feature space.<sup>141</sup> To support semantic labeling of RGB-D scenes (a database of common household objects recorded in 3D video), which is critical to many intelligent applications, a novel Long Short-Term Memorized Context Fusion model was developed to capture and fuse contextual information from multiple channels of photometric and depth data and were incorporated with CNN for end-to-end training for improved accuracy of fine-scale semantic labelling.<sup>142</sup> A number of other articles discuss data fusion and deep learning for driverless cars where multi-modal sensor information is fused using convolutional neural networks to better classify incoming information, often in real-time and space.<sup>143,144</sup> One study also used RNNs with Long Short-Term Memory to predict driver manoeuvres several seconds before they occur.<sup>145</sup> Finally, three articles focused on data fusion for improved diagnosis of disease.<sup>146-148</sup>

Interest in finding ways to integrate data fusion into data analytics and decision support is on the Pentagon's radar. On January 9, 2017, the Pentagon's Defense Innovation Advisory Board added an interim recommendation to their fact sheet for their public meeting that described "Establish(ing) a Global and Secure Repository for Data Collection, Sharing and Analysis". This recommendation includes the development of "tools to enable better decision-making through data mining, analysis and visualization" amongst other points. Eric Schmidt, executive chairman of Google, who is a member of the Defense Innovation Board, explained that this recommendation is based on the notion of data fusion: "The fantasy goes something like: we're going to have all these different signals; the signals will be automatically detected; the immediacy...will enable the warfighter to make a better decision".<sup>149</sup> Although this specific recommendation does not make reference to deep learning, other recommendations in the same document propose establishing a DoD center for the study of AI and ML (and presumably deep learning) in order to build expertise and capacity in the area with support from industry and academia suggesting that the data fusion and decision support would likely be facilitated by deep learning techniques.<sup>150</sup>

### Uncertainty analysis

The *Uncertainty analysis* group has 35 publications, 24 of which were published in 2015-2016. The majority of these recent articles discuss the issue of uncertainty modelling in deep learning based image recognition or speech recognition applications. One of the limitations of deep learning is that there is no measure of the uncertainty associated with their outputs. The consequence of this, particularly in terms of decision making, is that it makes it difficult to determine whether the data, model or decisions are reliable or are more akin to random guesses. There are two main types of uncertainty that are of concern in deep learning applications: *data uncertainty* and *model uncertainty*. *Data uncertainty* results from noisy data. *Model uncertainty* has two variations, resulting from uncertainty in model parameters (i.e. when a large number of models might explain a given dataset there is uncertainty about which model parameters to choose to predict with) or model structure uncertainty (i.e. uncertainty about which model structure to use). Both *data uncertainty* and *model uncertainty* can be used to attain predictive uncertainty, which provides a confidence level in the prediction or classification provided by the deep learning model.<sup>151</sup>

CNN's are particularly adept in computer vision and image recognition applications but have been shown to be over-confident in their predictions (i.e. overfitting). Various techniques for modelling uncertainty in deep neural networks have recently been proposed, including (but not limited to) dropout modelling, model averaging, and Gaussian processes. In one study, a Bayesian approximation of Gaussian processes was used to compute a predictive mean and variance which was then used to

determine confidence levels. The confidence interval was then used to improve the results of a CNN.<sup>152</sup> Good estimates of model uncertainty are also needed in 2D and 3D object recognition tasks where there is a large variability of shape cues. As such, one study found a stochastic gradient Markov Chain Monte Carlo method to be an effective technique to learn weight uncertainty in DNNs, resulting in higher recognition accuracy.<sup>153</sup> Another study proposed a weight uncertainty semi-restricted Boltzmann machine (WSRBM) to improve image recognition and image reconstruction, and was found to be more effective compared to the dropout method.<sup>154</sup>

Automatic speech recognition (ASR) suffers more from data uncertainty as its abilities degrade in the presence of noise and other distortions. Noisy speech is typically preprocessed using speech enhancement algorithms. However their accuracy is often inconsistent and varies widely across time-frames and can cause additional distortions in speech signals. The resulting uncertainty can be used through an uncertainty decoding method to account for the imperfections and improve ASRs.<sup>155,156</sup> Uncertainty decoding is a framework that uses the measures or estimates of uncertainty in the input features during acoustic model scoring. When combined with DNN-based ASR systems, various methods of uncertainty decoding, including Monte-Carlo sampling, Hidden Markov Models, and others, were shown to improve accuracy.<sup>156-160</sup>

### Risk analysis

*Risk analysis* is the second smallest group in the human decision support subset, with only 27 articles, 85% of which were written in 2015-2016. Following a review of all the articles, 15 were found to be related to health risks, while five were related to financial or credit risks. The health related articles mainly focus on predicting risk of particular diseases, such as breast cancer<sup>161</sup> or Alzheimer's disease,<sup>162</sup> but also include an article on predicting hospital readmission<sup>163</sup> and using electronic health record data to predict future diseases (more generally).<sup>164</sup> In terms of the credit risk, deep learning techniques were found to provide more accurate credit risk assessments<sup>165</sup> and sometimes perform better than standard feature selection methods that are currently being used by credit scoring agencies.<sup>166</sup> Deep learning methods were also found to be more useful than standard financial modelling techniques at designing and pricing securities, constructing portfolios and performing financial risk management by being able to detect and exploit interactions in the data that are generally undetectable by existing financial economic theories.<sup>167</sup>

Five other articles were classified into a group that is likely more relevant to defence interests as it covers issues related to security threats. In one article, CNN was used to detect the smuggling of non-explosive devices known as small metallic threats (SMTs) that have recently been used in terror attacks. CNN can detect SMTs in X-ray images of cargo containers more quickly and accurately than human visual inspection.<sup>89</sup> CNN has also been used to detect buildings from high-resolution satellite images for the purpose of (amongst others) disaster risk management in rural areas and villages.<sup>168</sup> One interesting study used deep learning technologies to bypass Google's reCaptcha service and Facebook's image captcha service<sup>a</sup> in an effort to identify reCaptcha's weaknesses and propose a series of modifications.<sup>169</sup> Another article reported on the use of deep learning to classify posts on a popular Chinese "end of the world" forum, in order to assess societal risk perception.<sup>170</sup> Finally, a joint Canadian-Chinese paper presented a methodology for critical infrastructure vulnerability assessment and risk analysis and identified the role of deep learning methods in the development and commissioning of future cyber defense solutions.<sup>171</sup>

---

<sup>a</sup> Captcha are text recognition or image based challenges designed to prevent robotic fraudsters from illicit activities online.

**Decision making/support (human)**

The *Decision making* group has 97 articles, 65 of which were published in 2015-2016. The largest portion of the articles in the last two years is related to medical applications and, in particular, those that involve image analysis for diagnostic purposes.<sup>172-174</sup> For example, one study describes a novel CNN to support neuroimaging-based brain disease diagnosis.<sup>116</sup> In a slightly different direction, but still in the medical applications domain, one article described the advantages and disadvantages of various deep-learning approaches for facilitating knowledge, pattern and relationship discovery in the growing publicly available unstructured medical research text corpora, which can be exploited by medical recommendation systems to support decision making.<sup>175</sup>

Another large group of articles was related to data mining and text analysis. Two focused on using sentiment analysis to support decision making. The first proposed a framework of sentiment analysis based on deep learning to deal with the high dimensionality of product reviews in an effort to support customer decision making.<sup>176</sup> The second tackled sentiment analysis of news articles and used the results to help semantically model the full text by using CNN with a bi-direction Long Short-Term Memory neural network to capture contextual information and extract salient information from the news documents.<sup>177</sup> Similarly, another study attempted to detect deceptive opinion spam in product and services reviews by using CNN to represent the semantic meaning of sentences in the review and to provide a document-level representation.<sup>178</sup> Another study combined deep learning with Decisional DNA which explores knowledge extraction based on a domain's past decisional events. Results of this study were found to be very promising in terms of enhancing the accuracy of knowledge-based predictions that are required in complex decision-making.<sup>179</sup> Finally, one study used stacked denoising auto-encoders to mine unstructured text data produced in cross-enterprise social interaction media, to map out massive manufacturing relationships which were used as decision support information.<sup>180</sup>

Some of the other articles published in the last two years looked at using deep learning to improve managerial decisions by:

- constructing a forecasting model based on customer behavior and sales data,<sup>181</sup>
- making decisions about credit applicants based on their credit scores,<sup>166</sup> and
- detecting and semantically tagging moving objects in video in near real-time speed to support intelligent planning and decision making.<sup>182</sup>

One article in particular presented an intelligent decision support tool that used a hybridized deep belief network and non-linear kernel-based parallel evolutionary support vector machine to predict and classify the evolutionary states of complex systems, such as early-warning-of-risk of harmful algal blooms and mobile robot navigation.<sup>183</sup>

Many of the remaining articles were not specifically focused on deep learning in decision making tools as much as deep learning for other applications that might include a step or component that involves or relates to decision making. For example, one article about using deep learning to mine data collected from the Internet of Things (IoT) mentioned that the data gathered by the IoT can be used to enable better decision making, improve efficiencies, and enhance productivity and accuracy.<sup>184</sup> Similarly, in an article about autonomous vehicles, the authors noted that this technology is about decision informatics and “embraces the advanced technologies of sensing (i.e., Big Data), processing (i.e., real-time analytics), reacting (i.e., real-time decision-making), and learning (i.e., deep learning)” but doesn't discuss the impact that deep learning has on decision making.<sup>178</sup> In these articles, and many others like them, the decision making is a sidebar to the focus of the article.

Recent news releases corroborate the growing trend of using deep learning to support decision making in the medical field. The Center for Digital Health Innovation at the University of California, San Francisco (UCSF) announced on January 19, 2017 a collaboration with Intel Corporation to build and test a medical based deep learning analytics platform designed to predict patient outcomes and support better front-line clinical treatment decisions. This next-generation platform is intended to integrate the huge volume and variety of data from genomic sequencing, monitoring, sensors and wearables to support rapid data-driven clinical decision making.<sup>185</sup> Transparency Market Research projects the global healthcare natural language processing (NLP) market to grow to US\$4.3B by the end of 2024, representing a CAGR of 18.8% between 2016-2024. NLP, based on machine learning and deep learning algorithms, is used in the medical field for predictive analysis and clinical decision support as the algorithms extract knowledge from clinical documents.<sup>186</sup> The U.S. FDA approved the first deep learning cloud-based medical imaging application, Arterys Cardio, to be used in clinical settings to diagnose heart problems.<sup>187</sup>

Deep learning decision support is also breaking into other industries such as human resource management and manufacturing, as well as at the CIA. Gayle Sheppard, who joined Intel corporation with the acquisition of the cognitive computing company Saffron Technology, projects that 2017 will be the year that AI-enabled decision support systems will move from testing to enterprise adoption across a variety of industries.<sup>188</sup> A December 2016 report by McKinsey & Company identified healthcare, smart cities, and human resource management as key industries to benefit from deep learning based decision support by making decisions more consistent, reliable and transparent.<sup>189</sup> Strategic Business Insights, another research and consulting company, also noted opportunities for deep learning in manufacturing in a diverse range of decision support tasks including production scheduling and sales strategy.<sup>190</sup> Finally, the Deputy Director for Digital Innovation at the CIA, Andrew Hallman, explained that the agency:

...has significantly improved its 'anticipatory intelligence' using a mesh of sophisticated algorithms and analytics against complex systems to better predict the flow of everything from illicit cash to extremists around the globe. Deep learning and other forms of machine learning can help analysts understand how seemingly disparate data sets might be linked or lend themselves to predicting future events with national security ramifications.<sup>191</sup>

Despite the progress that is increasingly being seen in the uptake of deep learning decision support, one challenge that remains is the commonly held view that deep learning is a black box from which it is hard to understand how results were reached. As such, users may continue to be reluctant to rely on deep learning outputs for decision making until the field is better able to explain what is going on under the hood.<sup>189</sup>

### **Military**

The *Military* category is the smallest in the dataset with only 22 publications, 68% published in the last two years. Six are related to some sort of target, threat or hazard detection and four are related to vehicle detection. Amongst the first group of articles, DBN was used to detect explosive hazards at range using forward-looking ground-penetrating radar (FLGPR) systems both with significantly increased probability of detection and nominal number of false alarms.<sup>192</sup> CNNs have been used in similar scenarios to classify GPR images of buried explosive hazards.<sup>193</sup> Research at the Pacific Northwest National Laboratory, under the auspices of the U.S. Army Public Health Command, has developed INSTINcT, an Intelligent Signature Canvas framework to detect depleted uranium aerosol doses and risks, using CNNs. CNNs have also been used for detecting typical targets from massive high-resolution remote

sensing image data.<sup>194</sup> In terms of vehicle detection, research has focused on both stationary and moving vehicle detection, including aircraft detection, from Laser Doppler Vibrometry,<sup>195</sup> UAVs (more generally),<sup>196</sup> satellite imagery,<sup>197</sup> and very high resolution remote sensing images.<sup>198</sup> The addition of deep learning algorithms helps overcome challenges with low-quality UAV videos and platform movement (i.e. vibrations), and makes significant contributions including real-time and higher detection accuracy rates. Identification of different vehicle types can help the military recognize and distinguish details, such as those with missile launchers or mounted machine guns.<sup>199</sup>

Two other articles of interest looked at the use of deep learning for situation awareness. The first is a 2016 publication by the National Defense University, China, which analyzes DARPA's use of deep learning technologies for battlefield situation intelligent cognition.<sup>200</sup> The second is a 2015 Pennsylvania State University dissertation on dynamic data-driven application systems (DDDAS), which integrate simulation with dynamically assimilated data, multiscale modelling and computation that combines symbolic time series analysis and deep learning. DDDAS have potential applications for situation awareness in intelligence, surveillance and reconnaissance missions.<sup>61</sup>

Overall, the use of deep learning in military applications is not frequently presented in the academic literature, but this is not likely to remain the case as it is known from a few searches on military web information resources that the military is actively pursuing this area. For example, DTIC has roughly 111 unique entries from 2011-2016 that touch on deep learning, 34 of which were from the last two years. Overall, these documents, although not analyzed as part of this project,<sup>b</sup> are more focused on image analysis followed by acoustic modelling, speech recognition and computer vision.

A brief selection of recent U.S. defense and intelligence research in the area of deep learning also reflects a continued move in this direction. In 2016, AFRL-Rome released broad agency announcements related to deep learning including:

- “Deep Learning for Actionable Intelligence Discovery and Exploitation”: a US\$9.9m business opportunity seeking the development and demonstration of a time-dominant data fusion and closed loop ISR system for intelligence customers, and
- “Multi-Source Information Extraction and Network Analysis (MUSIENA)”: a call for white papers on developing MUSIENA capabilities for the Air Force to better conduct analytical operations in support of ISR missions.<sup>201</sup>

Additionally, IARPA's Creation of Operationally Realistic 3D Environment (CORE3D)<sup>202</sup> and Deep Intermodal Video Analytics (DIVA)<sup>203</sup> research programs include research on deep learning. In August 2016, DARPA released a broad agency announcement called “Explainable Artificial Intelligence (XAI)” soliciting innovative research proposals to create a suite of new or modified machine learning techniques, including deep learning, that produce explainable models that end users can easily understand.<sup>204</sup> DARPA's 2017 President's Budget Submission includes mention of deep learning in a number of projects:

- Mining and Understanding Software Enclaves(MUSE): “Apply deep learning algorithms on complex graph structures produced by corpus mining to discover latent relationships among corpus elements for automated program repair and synthesis.”

---

<sup>b</sup> It is not possible to export the standard metadata from DTIC for that is used in our analyses. As such, titles and document descriptions were imported into Vantage Point and natural language processing was performed on the text. The keywords were then grouped using the existing thesauri from this project, and the results are reported here at a high level.

- Automatic Target Recognition (ATR) Technology: “Develop adaptable ATR algorithms to rapidly learn new targets with minimal measured data and evaluate algorithm learning rate.”
- Seeker Cost Transformation (SECTR): “The technical approach to target recognition will start from “deep learning” and 2D/3D machine vision algorithms pioneered for facial recognition and the identification of critical image features.”<sup>205</sup>

In a 2015 market report, *Frost and Sullivan* noted that the rising use of UAVs to provide ISR data to the military creates a growing demand for deep learning technologies to provide automated analysis of the mounting petabytes of data that is captured. Deep learning technologies are projected to assist analysts in dealing with the continual flood of information.<sup>206</sup> Finally, general web searching revealed that MIT, funded by DARPA, has created a new microchip that is intended to perform deep learning functions on smart phone-sized devices in the battlefield,<sup>207</sup> and the Pentagon is planning on using deep learning to better understand ISIS as a network and how to better target and defeat it.<sup>208</sup> Given that there is clear interest by the military in using deep learning, it is likely that more academic research, directly tied to military interests, will be forthcoming.

### 5.3 Quantum Computing

Overall there is limited evidence in the dataset that quantum computing will have a significant impact on deep learning. The full dataset contained only 16 articles (0.002%) that included the keyword “quantum”, only five of which showed up in the sub dataset. One article based on research conducted in 2014 by researchers at Microsoft showed that quantum computations provided a much better framework than classical approaches (e.g. contrastive divergence) to identifying algorithm weights and training Boltzmann Machines. The authors note, however, that while numerical results on small examples are encouraging, scalable quantum computers will be needed to assess the generalization of their findings.<sup>209</sup> Researchers at Lockheed Martin ran similar tests and found that quantum sampling-based training of a Restricted Boltzmann Machine achieved comparable or better accuracy with less iteration than conventional contrastive divergence training. They likewise cautioned that further investigations were needed to determine to what extent the improvements could be attributed to quantum effects.<sup>210</sup>

A few companies do appear to be attempting to use quantum computers to support deep learning applications. D-Wave (a Canadian company, the first in the world to sell quantum computers) is working on exploiting the parallels between sparsely connected restricted Boltzmann machines and quantum computers to substantially accelerate learning in deep, hierarchical neural networks.<sup>211</sup>

At the same time, Google explained in 2015 that it was not using the quantum computer it acquired from D-Wave to support its deep learning applications, as it viewed the quantum computing infrastructure as not being the best fit for systems such as CNN or RNN. One of the factors that lead to this conclusion was the fact that deep learning requires both a model and a set of values for parameters, while the number of parameters and operations that current quantum computers can hold are very small.<sup>212</sup>

It appears that more time and research (by the deep learning and quantum computing communities) will be needed to determine if, or in what scenarios, there is a viable match between the two.

## 6 KEY PLAYERS

### 6.1 Key Players: Full Dataset

Figure 9 shows the top 20 author affiliations. The majority of the top affiliations are from China, (8), followed by USA (5), and Singapore (4). Three major corporate bodies are near the top of the list, including Microsoft (284 publications), Google (141) and IBM (127) each with international corporate and research locations. Both Google and IBM include publications from notable subsidiaries or divisions in the field, including Google’s DeepMind and the IBM T.J. Watson Research Center, where the majority of IBM’s research is from. Three Government/RTO affiliations in this list include the Chinese Academy of Sciences, which is the top affiliation with 445 publications, the Agency for Science, Technology and Research (A-STAR), Singapore (82) and the Centre National de la Recherche Scientifique (CNRS), France (73). One Canadian affiliation, the University of Toronto (85), appears in the 15<sup>th</sup> position in this list.

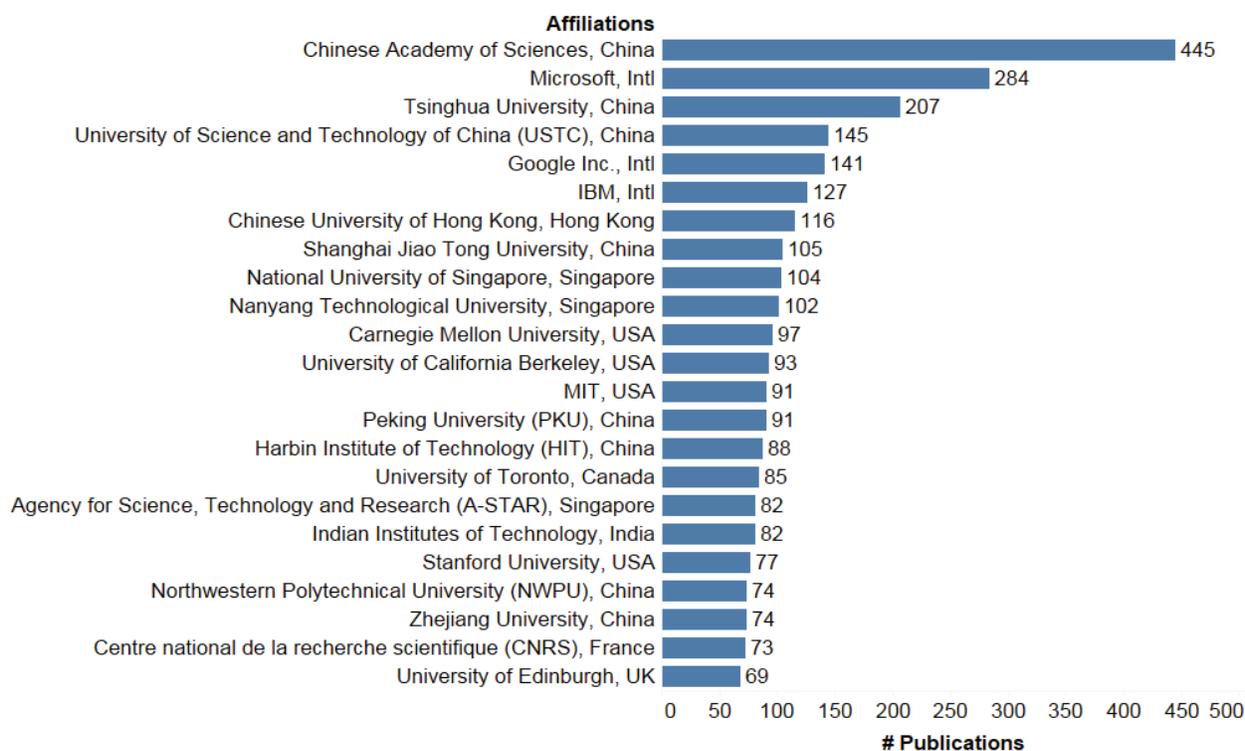


Figure 9. Top Affiliations, Full Set

Canada is actually one of the top five countries overall in the dataset with a total of 403 publications, preceded only by USA (2779), China (2601), UK (618) and Japan (408). Figure 10 presents the top Canadian affiliations with 10 or more publications.

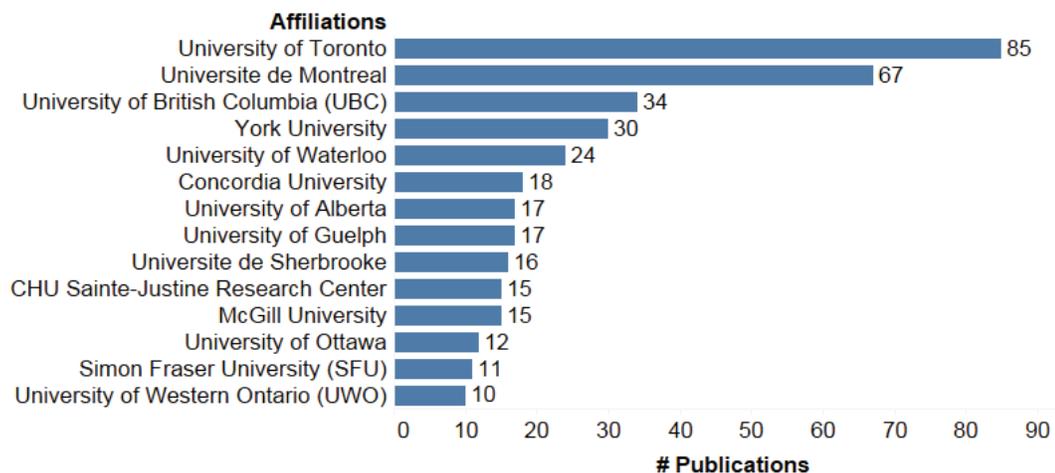


Figure 10. Top Canadian Affiliations, Full Set

### 6.1.1 Canadian Areas of Expertise

The major areas of expertise (five or more publications in the full dataset) for the top Canadian players are presented in Figure 11. A full list of research areas is available in the Tableau workbook. Canada appears to be more broadly focused on *Image analysis & processing* than *Speech recognition*.



Figure 11. Areas of Research, Canadian Affiliations, Full Set

### 6.1.2 International Collaborations

All collaboration maps in this report were produced using TouchGraph Navigator<sup>c</sup> software. Nodes within clusters tend to be close together while nodes in separate clusters are farther apart. That being said, the proximity of nodes has been slightly altered to improve the readability of the map.

<sup>c</sup> TouchGraph Navigator is produced by the US company TouchGraph LLC: <http://www.touchgraph.com/navigator>

TouchGraph’s clustering algorithm ignores long edges that connect separate clusters, thereby splitting separate connected components into the clusters that are shown.<sup>d</sup>

Figure 12 presents the top 20 affiliations, as well as Canadian affiliations with a minimum of three publications and their collaborations with all other affiliations in the full dataset that have a minimum of ten publications. This figure, which shows a minimum of four collaborations, along with another collaboration map that is filtered to seven collaborations, are available in the Tableau workbook. There are six major clusters in this map. The main two are centred on the Chinese Academy of Sciences (turquoise) and Microsoft (purple), with many shared collaborative partnerships between the two. Another two are centered on Google (yellow) and IBM (dark blue), both of which predominantly include North American collaborators. The last two clusters involve CNRS, France (dark green, bottom) and Université de Montreal (bright green, top) both of which are dominated by collaborators in their home countries. Canadian players are circled on the map.

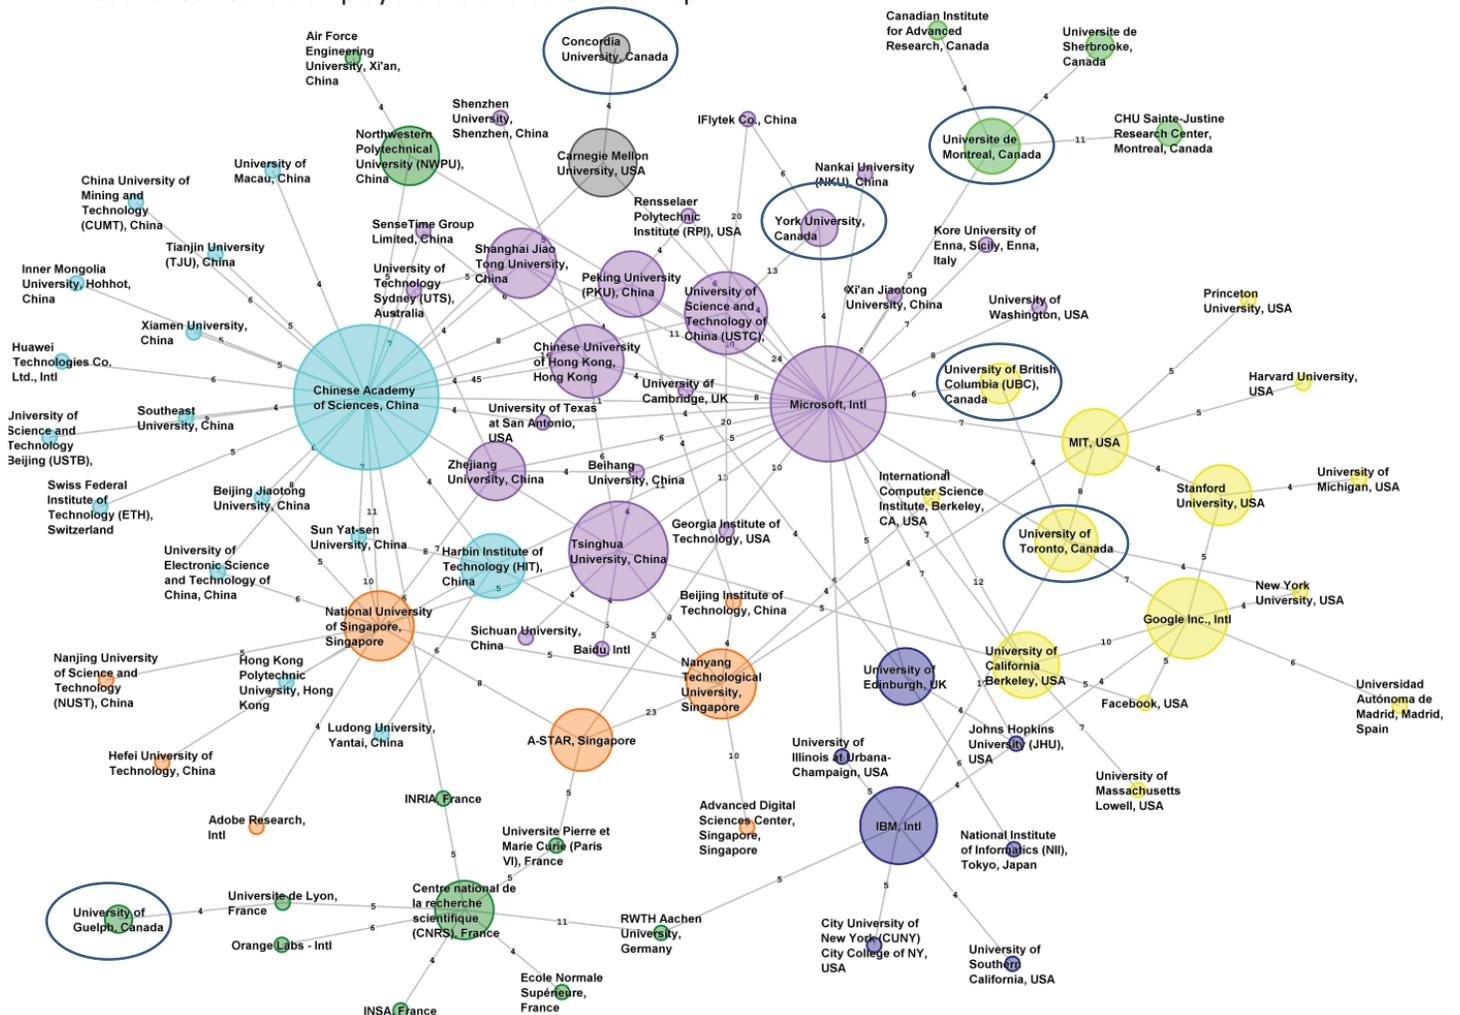


Figure 12. International Collaborations, Full Set

<sup>d</sup> TouchGraph Navigator identifies clusters based on ‘Edge Betweenness Centrality’. For more information on the clustering algorithm used by TouchGraph Navigator please see [http://www.touchgraph.com/assets/navigator/help2/module\\_7\\_7.html?MenuState=AUDAAO-0hAHvtIQABO-0hO-0gWAZAQPvtIRjBHvZdGvY](http://www.touchgraph.com/assets/navigator/help2/module_7_7.html?MenuState=AUDAAO-0hAHvtIQABO-0hO-0gWAZAQPvtIRjBHvZdGvY)

### 6.1.3 Canadian-International Collaborations

Figure 13 presents the international collaborations of the Canadian affiliations that have a minimum of three publications in the full dataset. The map is filtered to show collaborations that include a minimum of two co-publications. The map reveals a central cluster surrounding the University of Toronto (U of T), Université de Montreal (U de M), York University, and the University of British Columbia (UBC) which have all collaborated with Microsoft at least four times. U of T and U de M have also both collaborated with Facebook, the Canadian Institute for Advanced Research (CIFAR), Université de Sherbrooke and New York University but never directly with each other. Each of the four main Canadian players in this cluster has many other international collaborative partnerships. U of T has the highest number of repeat collaborations on the map with 10 co-publications with IBM. Additional clusters include one at the University of Waterloo with a number of Chinese and European affiliations, and three other smaller collaborative relationships for Concordia University, Dalhousie University and Simon Fraser University.

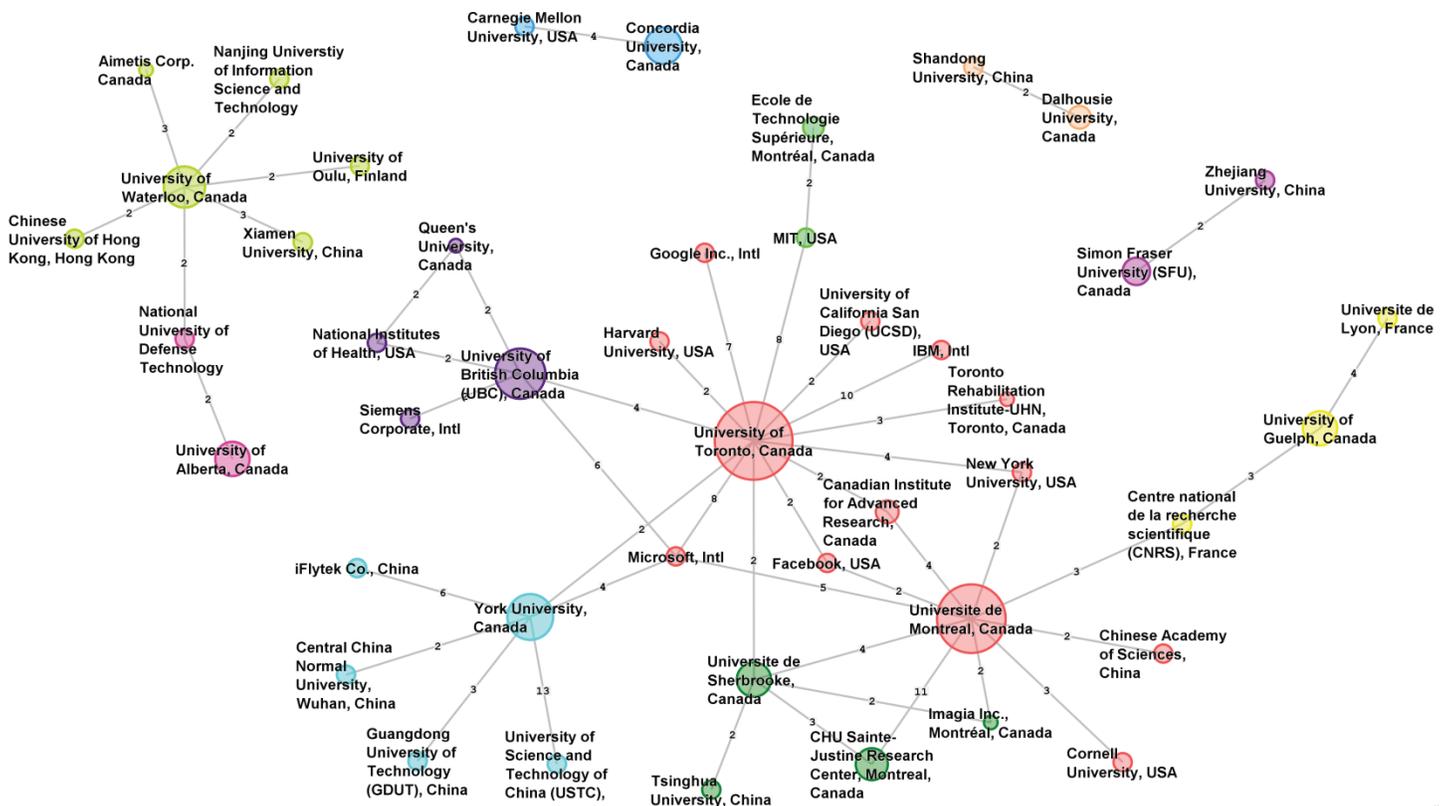


Figure 13. Canadian-International Collaborations, Full Set

### 6.2 Key Players: Human Decision Support Subset

Figure 14 presents the top affiliations in the human decision support subset. While many of the same affiliations from the full dataset are present, a few new affiliations appear, including: Beihang University, China (25), University of Electric Science and Technology of China (23), University of North Carolina, USA, University of Oxford, UK and Xiamen University, China, all with 22 publications.

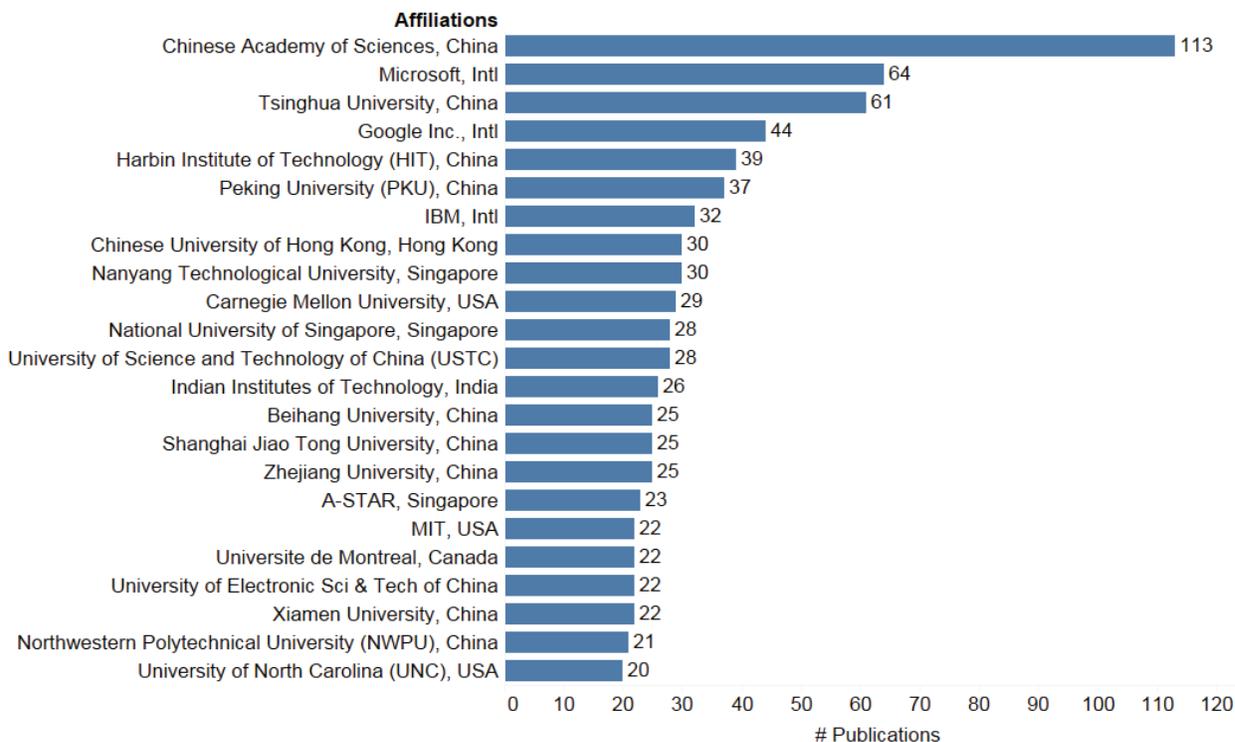


Figure 14. Top Affiliations, Human Decision Support Subset

Figure 3. Subjects in Human Decision Support Subset  
 Figure 15 presents the top Canadian affiliations in the human decision support subset with a minimum of three publications. New to this list are CIFAR with four publications, and the University of Alberta with three publications

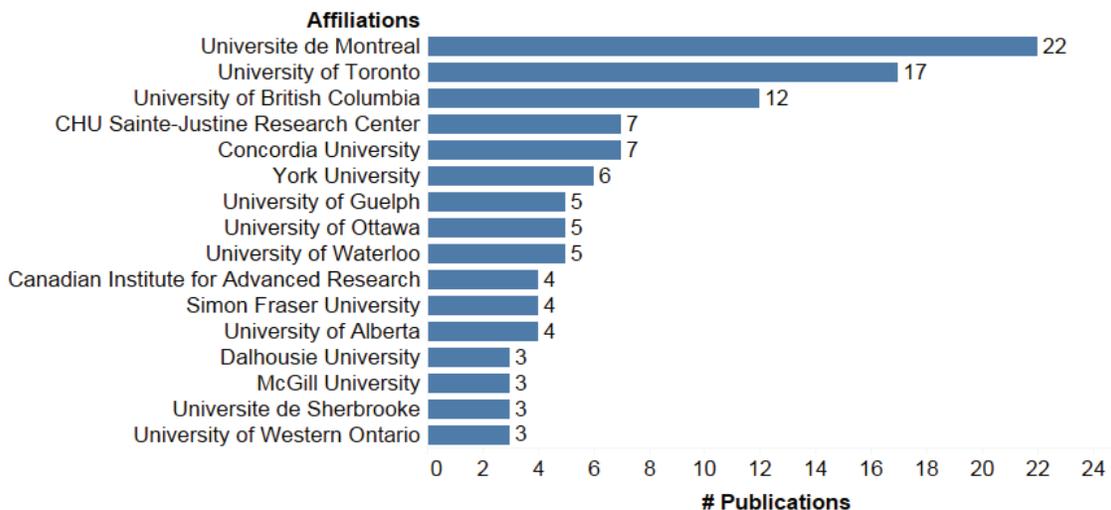


Figure 15. Top Canadian Affiliations, Human Decision Support Subset

### 6.2.1 Canadian Areas of Expertise

Figure 3. Subjects in Human Decision Support Subset Figure 16 presents the research focus of the Canadian affiliations in the areas of human decision support. The figure is limited to those topics with a minimum of three publications per affiliation. A summary of each affiliation’s research in these areas is provided below.

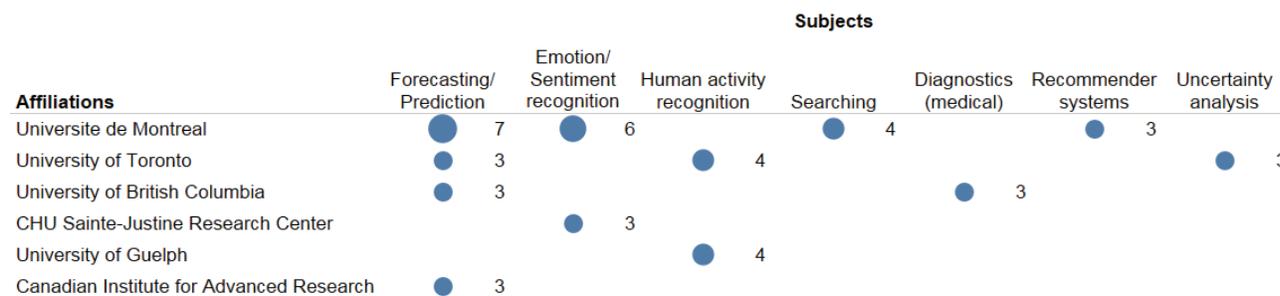


Figure 16. Canadian Affiliations Areas of Expertise

The Université de Montreal (U de M) has 22 publications in the decision making subset and 67 publications in the full dataset. U de M is the academic home of Yoshua Bengio, one of the pioneers of the modern stream of deep learning research. Bengio is currently head of the Montreal Institute for Learning Algorithms (MILA) in the Department of Computer Science and Operations Research, as well as the co-director of the CIFAR program on Learning in Machines and Brains (formerly known as Neural Computation & Adaptive Perception). He is also the Canada Research Chair in Statistical Learning Algorithms.<sup>213</sup> MILA is world-renowned for many breakthroughs in developing novel deep learning algorithms and applying them to language modelling, machine translation, object recognition, structured output generative modelling and speech recognition.<sup>214</sup>

Publications from the U de M on *Forecasting/Prediction* (seven articles) mainly focus on testing or developing new or improved deep learning algorithms with improved ability to predict accurately. For example, one article investigates the modelling of time series data using gated auto-encoders and recurrent networks, which outperform standard recurrent networks in terms of prediction accuracy on a variety of tasks.<sup>215</sup> A paper co-published with University of British Columbia, reveals that many deep learning models can accurately predict the parameter values (or weights of a network) with minimal to no training.<sup>216</sup> Another publication introduces the multi-prediction deep Boltzmann machine (MPDBM) which, unlike the traditional DBM, does not require greedy layerwise pre-training and outperforms DBM at classification, classification with missing inputs, and mean field prediction tasks.<sup>217</sup> From a different perspective, another article presents a deep learning approach to predicting velocity and direction changes (odometry) from visual information using CNN.<sup>218</sup>

Six of U de M’s articles are focused on *Emotion/Sentiment recognition* with three being co-published with the CHU Sainte-Justine Research Center. These three co-publications are focused on emotion recognition in video and emanated from the U de M’s winning team submission to the 2013 Emotion Recognition in the Wild Challenge. The articles report on a variety of approaches including the use of a multi-modality classifier (ie CNN, DBN, K-means and auto-encoder)<sup>219,220</sup> as well as a hybrid CNN-RNN architecture.<sup>221</sup> Other studies investigate deep learning to perform emotion recognition in 2D face images,<sup>222</sup> in online reviews and recommendations,<sup>223</sup> and in modelling affect manifested through physiological signals of video game players.<sup>224</sup>

Four of U de M's articles are focused on the use of deep learning in *Searching*, two of which focus on the development of latent semantic models using deep learning,<sup>225,226</sup> and one that identifies random search as more effective for hyper-parameter optimization in neural networks and deep belief networks than grid search.<sup>227</sup>

One of the more interesting articles published by the U de M with the Harbin Institute of Technology, China on recommender systems tested a CNN-based model which used social media data to identify users' consumption intentions. The new model (Consumption Intention Mining Model – CIMM) was found to be both effective and transferrable across domains (which is typically challenging with other CNN models) and is thus a powerful tool for recommender systems and targeted advertising.<sup>228</sup>

The University of Toronto (U of T) has 17 publications in the decision making subset and 85 in the full dataset. The U of T Computer Science department is home to Geoffrey Hinton, another pioneer of the modern stream of deep learning research, who also works part time at Google as an engineering fellow.<sup>229</sup> Hinton is a member of the Machine Learning research group which focuses on neural networks, statistical pattern recognition, probabilistic planning and adaptive systems, as well as a member of the Toronto Deep learning group whose [website](#) provides live demos of image classification and image to text conversion.

Four of U of T's publications are focused on human activity recognition, each of which presents a new deep learning architecture to perform one or more of a variety of human activity recognition tasks. One of the articles proposes a new feed-forward, auto-regressive architecture to perform facial expression recognition.<sup>230</sup> Two examine a combined hierarchical Dirichlet process Deep Boltzmann Machine (HDP-DBM) model to learn novel concepts from very few training examples that could efficiently learn and make inferences in object recognition, handwritten character recognition and human motion capture datasets.<sup>231,232</sup> The last describes the development of a Deep Belief Network (DBN) which could be used for recognition tasks, including gesture, object and facial expression recognition.<sup>233</sup>

U of T's three publications on *Forecasting/Prediction* are, similar to U de M, focused on the predictability of the deep learning architectures that are being tested. For example, one article, published with Google, describes using an autoregressive product of a DNN-Hidden Markov Model system trained to predict the phone labels of multiple frames to improve speech recognition accuracy.<sup>234</sup> Another discusses the challenge of overfitting in large networks and the ineffectiveness of using predictions of many different large neural nets to overcome this challenge.<sup>235</sup> The last, co-published with CIFAR, reports on the use of deep learning as an approach to pattern discovery in predicting the sequence specificities of DNA- and RNA-binding proteins.<sup>236</sup>

U of T's three publications on *Uncertainty analysis* focus on using deep learning to reduce uncertainty involved in tasks. For example, the authors of one paper used CNN to perform flow matching (of traffic participant motion in the context of autonomous driving) and showed that the CNN was able to estimate the uncertainty of its matches.<sup>237</sup> Another article investigates the use of deep (factored)-restricted Boltzmann machines for simultaneous object tracking and recognition that includes an attentional mechanism that learned to control gazes so as to minimize tracking uncertainty. The last article describes the use of Bayesian CNN that can infer sun direction in images as part of a visual odometry pipeline that computes a principled uncertainty associated with each prediction which is then used for optimal data fusion.<sup>238</sup>

The University of British Columbia (UBC) has 12 publications in the decision making subset and 34 in the full set. UBC's Department of Computer Science has an artificial intelligence unit with a machine learning research group lead by Mark Schmidt that focuses on, amongst other topics, deep learning, neural computation and prediction, and optimal decision making.<sup>239</sup>

Half of the articles in the subset are focused on either *Forecasting/Prediction* (3) or *Diagnostics (medical)* (3) research. The first of the *Forecasting/Prediction* articles describes the use of deep learning to extract latent features from brain lesion patterns to predict short-term disease activity in patients with early MS symptoms.<sup>113</sup> A second article shows that algorithms using deep stacking networks are more effective than existing multiple measurement vectors for modelling compressive sensing settings as it improves the shallow neural network prediction.<sup>240</sup> Last is a co-publication with U de M which reveals that many deep learning models can accurately predict the parameter values (or weights of a network) with minimal to no training.<sup>216</sup> The three articles on *Diagnostics (medical)* report on various uses of deep learning for medical image based diagnoses including :

- the use of deep belief networks as part of a framework to analyze temporal ultrasound signals of prostate tissue to complement and improve differentiation of benign and cancerous tissues in MRIs,<sup>241</sup>
- the use of deep feed-forward neural networks to automatically detect and predict the location of vertebrae in CT scans to improve computer-aided diagnosis systems,<sup>242</sup> and finally,
- the use of an improved convolutional deep belief network which is significantly faster than previous versions and makes training on 3D medical images practical.<sup>243</sup>

The Canadian Institute for Advanced Research (CIFAR) has only four publications in the sub-dataset and eight in the full dataset. It is the home of the Learning in Machines & Brains (LMB) research program (formerly known as Neural Computation & Adaptive Perception) which, according to Yann LeCun, "had a huge impact in forming a community around deep learning". LeCun is the third key pioneer of the modern stream of deep learning research<sup>e</sup> and co-director of the LMB program with Bengio (U de M). Hinton (U of T) acts as an advisor of the LMB program after formerly acting as director.<sup>244,245</sup> The low publication count might be the result of the fact that all of the research fellows are associated with other universities and thus likely publishing under their academic affiliation instead of under CIFAR.

Three of CIFAR's publications are found in the *Forecasting/Prediction* group. One recent article presents a Recombinator Networks model combined with a denoising prediction model that reduces error from the previous state-of-the-art on two facial keypoint datasets.<sup>246</sup> A second article discusses learning classification labels by using CNN to make predictions at the instance level by inferring sentence similarity.<sup>247</sup> The last article, co-published with U of T, reports on the use of deep learning as an approach to pattern discovery in predicting the sequence specificities of DNA- and RNA-binding proteins.<sup>236</sup>

The University of Guelph has five publications in the subset and 17 in the full dataset. The Guelph School of Engineering has a Machine Learning Research Group led by Graham Taylor, who studied under both G. Hinton and Y. LeCun. Taylor is also the CIFAR Azrieli Global Scholar in the LMB research program. Four of the University of Guelph's publications are on *Human activity recognition* and are mostly focused on

---

<sup>e</sup> Yann LeCun later headed up the AI research lab at Facebook and founded the NYU Centre for Data Science. (see <http://yann.lecun.com/> for more details on LeCun's work)

gesture recognition and hand motion. Three of the articles describe the value of using multi-scale and multiple data modalities (e.g. depth video, articulated poses and speech) along with a RNN to model large-scale temporal dependencies, data fusion and ultimately gesture classification. This approach was further improved by the use of a ModDrop training technique that lead to a robust classifier that could provide meaningful predictions when signals were missing from one or several modalities.<sup>248 249,250</sup> The remaining article proposes a deep learning based-approach for hand pose estimation, targeting and gesture recognition, that requires very little labelled data and uses unlabeled data during training to improve results.<sup>251</sup>

### 6.2.2 International Collaborations

Figure 17 presents the international collaborations in the human decision support subset limited to those with three or more occurrences. The three Canadian collaborators are encircled. Once again, the two main clusters in the map are centered on the Chinese Academy of Sciences, China and Microsoft. Google (upper left, red) and IBM (middle right, yellow) have only one collaboration each in this map while the CNRS, France is no longer present. U de M forms a cluster with CHU Sainte-Justine Research Center and Goethe University, Germany while the University of Guelph is seen collaborating with the Université de Lyon, France.



### 6.2.3 Canadian-International Collaborations

Figure 18 presents the collaboration networks of the Canadian affiliations with a minimum of 10 publications in the human decision support subset. The map is filtered to show collaborations that include a minimum of two co-publications. This map is significantly less populated although CNRS, France does reappear with connections to University of Guelph. Both U de M and UBC are collaborating with Microsoft, while UBC is collaborating with U of T, who is collaborating with Google and MIT, USA.

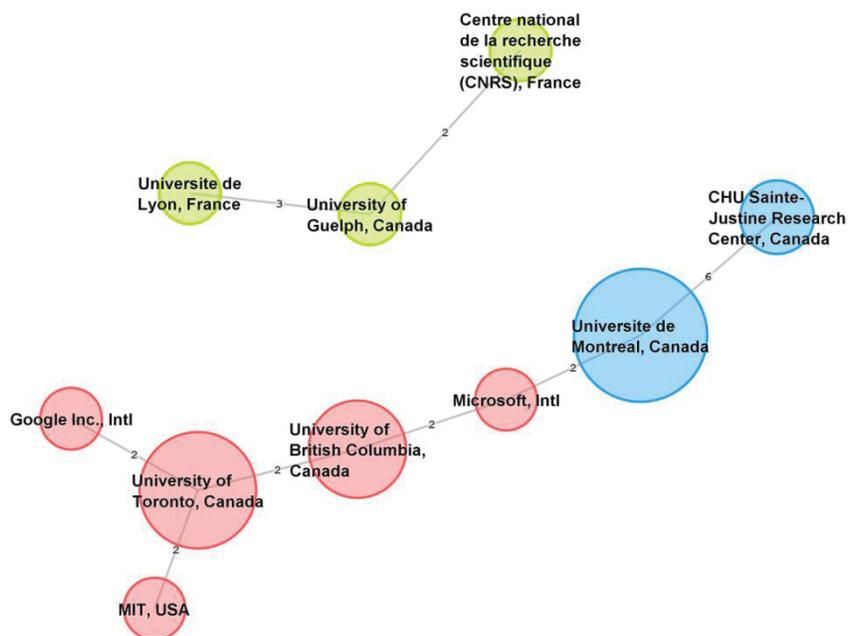


Figure 18. Canadian-International Collaborations, Subset

## 7 APPLICATION AREAS

Figure 19 presents a list of all applications from the full dataset that are also found in the human decision support subset.

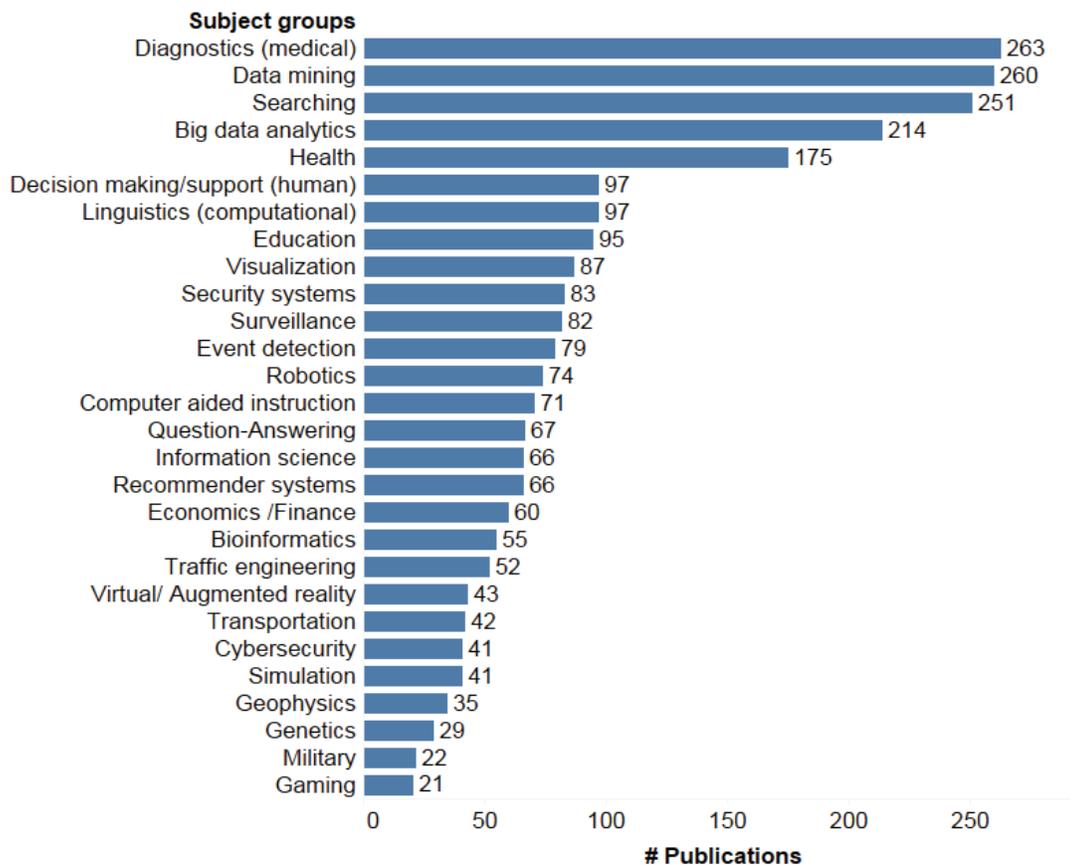


Figure 19. Applications in Subset

While this list is largely comprised of the preselected application-based subject groups that are in the human decision support subset, we also see the following applications:

- Bioinformatics
- Cybersecurity
- Economics/Finance
- Gaming
- Genetics
- Geophysics
- Health
- Information science
- Linguistics
- Robotics
- Simulation
- Traffic engineering
- Transportation

- Virtual/Augmented reality

This suggests that each of these application areas has some overlap with human decision making concepts. To identify the overlap, Figure 20 presents a co-occurrence matrix of the applications listed above with the human decision support subset topics. To facilitate viewing, the figure has been limited to co-occurrences of five or greater (the full matrix is available in the Tableau workbook). This figure allows us to see that, for example, deep learning research in *Decision making/support (human)* has occurred in the context of *Health, Robotics, Economics/Finance* and even *Gaming*. Similarly, deep learning research in *Data Mining* is relevant in *Health, Linguistics (computational), Economics/Finance, Information science, Bioinformatics* and *Geophysics*.

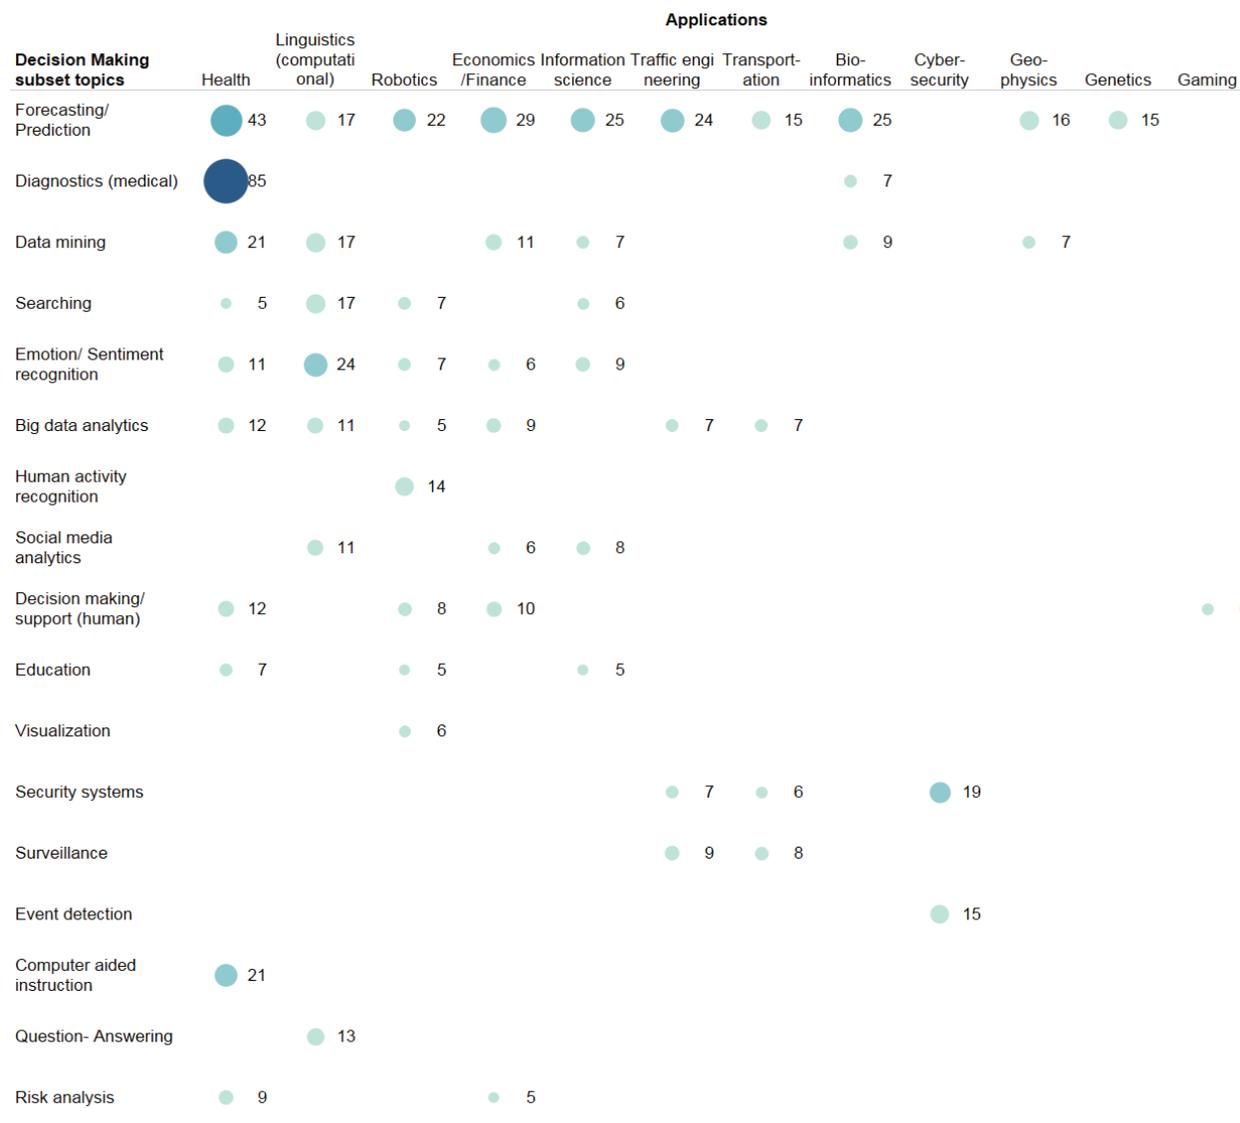


Figure 20. Co-Occurrence of Decision Support Topics with Application Topics

## 7.1 Future Applications

One of the challenges of identifying future applications is that the majority of the research is focused on advancing existing algorithms and reporting on their results, and is not theorizing about future applications or directions of research at this stage. The human decision support dataset has only 16 articles that explicitly mentioned future directions, and of these, particularly for those articles written before 2015, the goals have already been reached. Other sources that do talk about future goals in regards to deep learning are more focused on the larger AI question regarding the potential for AI to surpass human intelligence and what consequences that might have for society. Virtually no material was found discussing the future applications of deep learning as applied to human decision support.

Frost & Sullivan projected in 2015 that the huge amount of data generated by aerospace and defence, education, finance and healthcare made them key sectors to be impacted by deep learning. Deep learning was expected to be applied in these sectors to unravel customer insights, buying behavior, and future impact scenarios to support key decision makers. For example, deep learning of big data is expected to generate visualizations of massive datasets involved in portfolio and investment management supporting major strategic decision making. Similarly, marketing can make use of deep learning to reveal clear patterns of collaboration between developers and vendors to support strategic decision making.<sup>252</sup>

LeCun (Director of AI, Facebook) noted, at a 2016 CERN Colloquium called Deep Learning and the Future of AI, that while the most impressive applications of deep learning are emerging from industry, the majority of good ideas will continue to come from academia. In particular, he noted that the image and speech recognition of today will soon lead to better language understanding, dialog and translation. Similarly, today's image recognition capabilities at Facebook, Google, Twitter, Microsoft and others, will soon lead to better auto-pilots for cars, robot perception and medical image analysis. New hardware for embedded applications is still needed for smart cameras and mobile devices, to name a few. Finally, he explained that building deep learning applications that integrate reasoning, have a good architecture for short term memory, and can follow good principles for unsupervised learning, are still far off on the horizon.<sup>253</sup>

In a series of recent interviews with some of the top names in deep learning (e.g. Director of SkyMind, research scientists at OpenAI, Google, CIFAR, etc... ), the most frequently mentioned area of research that was expected to be advanced in the coming years was related to deep unsupervised learning. Other applications that were identified as likely in the next five years include:

- Automation of some decisions that are taken based on new knowledge gained from deep learning and AI based big data analytic tools. (Based on reinforcement learning and interpretation of sensory data)<sup>254</sup>
- Bridging the gap between big data and better decisions (by combining deep learning for perceptual tasks with other AI and machine learning techniques for reasoning)
- Models that can learn from fewer training cases
- Applications with multi-modal problems with more structured data, opening opportunities in data mining and knowledge discovery
- Applications that can run on cheap mobile devices without extra hardware support or prohibitive memory overhead

- Scaling to video; outperforming current approaches to natural language processing; significant advances in deep reinforcement learning
- Modelling the behaviour of genes, drugs, proteins for the design of new medicine
- Learning that persists beyond individual datasets
- Video recognition, medical imaging and text processing
- Integration of deep learning in commercial products (like face detection which has been incorporated in consumer cameras in the past 10 years).<sup>255</sup>

In a review of deep learning applications that are set to impress the market in 2017, the following applications were noted:

- Advanced melanoma screening and detection
- Brain cancer detection
- Ultrasound images, pre-natal care
- Weather forecasting and event detection
- Emerging market price forecasting
- Space mission efforts
- Civil and mechanical engineering – predict structural response under earthquake load.<sup>256</sup>

In a presentation at the 2016 GPU Tech Conference, B. Zhang, a technology fellow with BAE Systems, projected that future geospatial intelligence systems will incorporate deep learning and offer smarter systems that could learn from their mistakes, and detect and monitor defense relevant objects at 99% accuracy, with significantly reduced software engineering and enhancement costs.<sup>257</sup>

Finally, although not much information was found on future applications of deep learning in decision support applications, one indication that this market is expected to grow can be seen in the figure below regarding IBM’s “Cognitive Solutions,” which is based on Watson, which uses deep learning. Figure 21, which was presented at an IBM Investor Briefing in 2016 and projects that “Cognitive Solutions”, which support decision making, will be a ~\$2T USD market by 2025.<sup>258</sup>



Figure 21. IBM’s Projected Decision Support Market Opportunity

## 8 DEEP LEARNING LIMITATIONS

Given the current stage of deep learning research, which is being led by major corporations and is primarily focused on development and advancement of existing algorithms, there is currently very little documentation of deep learning limitations, especially within the specific context of human decision making support.

In general, there are three main limitations to deep learning, including:

- Architectures and output models that are, in many ways, seen as a black box,
- Deep learning's inability to reason or apply learning to another context,
- The need for massive amounts of data.

In terms of deep learning being a black box, this is a commonly used term to discuss the fact that deep learning suffers from a lack of theory supporting its processes, algorithms and resulting models. For example, given a certain dataset and goal, there could be two neural networks with different weights that achieve the same results, but how they have reached the results remains unexplainable. Similarly, while it might reach a decision with better accuracy than another technique, it cannot tell you why it reached that decision. While it is possible to interrogate the hidden layers to see what the machines are doing, an empirical confirmation of the outputs is generally needed. This is especially true given the many layers and thousands of nodes that may be at play. As such, data scientists are left with evaluating a model based on a measure of how well they predicted a result, rather than understanding the architecture itself. This leaves open the potential for blind spots that might not be caught by initial training and test data and might lead to unexpected errors in unusual situations.<sup>259, 260-262</sup>

Developing common sense reasoning or even "understanding" is a big challenge for all machine learning systems, including deep learning. Most deep learning algorithms are only good at a single task (no single algorithm is yet efficient at both identifying objects and playing video games) and they cannot transfer what they have learned from one domain to another. For example, Google's DeepMind Deep Q Reinforcement algorithm, which learned to play Atari games, could not transfer its learning from one game to another and had to learn each game from scratch, even if it had played it before. From a military perspective, while deep learning can rapidly identify objects and detect threats in massive amounts of ISR data, it cannot provide actionable insights based on the data input. While solutions to this specific example, and the problem in general, are in the works, they are being developed at a slower pace than the more technical abilities of, for example, image or text recognition.<sup>263,259,261,262,264</sup>

Last is the need for massive amounts of data to effectively train and test deep learning algorithms. Finding ways for deep learning tools to more efficiently work with less data is necessary for a wider range of industries to make use of it.<sup>261,265</sup>

Many human decision support tools will use big data as inputs, but the nature of big data brings challenges in terms of its: "high dimensionality, streaming data analysis, scalability of deep learning models, improved formulation of data abstractions, distributed computing, semantic indexing, data tagging, information retrieval, criteria for extracting good data representations, and domain adaptation".<sup>266</sup> In this context, the following questions were noted in 2015, as remaining to be answered:

- What volume of input data is generally necessary to train useful (good) data representations by Deep Learning algorithms, which can then be generalized for new data in the specific big data application domain?
- Domain adaptation which considers the distribution shift between the input data source (for training the representations) and the target data source (for generalizing the representations) is an important focus of study. This relates to the generalization capacity of deep learning data representations and patterns, which is an important requirement in big data analytics.
- What criteria are necessary and should be defined for allowing the extracted data representations to provide useful semantic meaning to big data (e.g. something akin to misclassification errors)? In other words, what would constitute a good data representation that is effective for semantic indexing and or data tagging?<sup>266</sup>

Finally, in a plenary speech at the IEEE-ICASSP meeting that was given by Li Deng, Chief Scientist of AI, Microsoft, in March 2016, the following challenges for future research were identified and echo some of the limitations discussed above:

1. Structured embedding for better reasoning: integrate symbolic/neural representations
2. Integrate deep discriminative and generative/Bayesian models
3. Deep unsupervised learning.<sup>267</sup>

This last point, the need to develop unsupervised deep learning, was echoed by LeCun (Director of AI, Facebook) at a 2016 CERN Colloquium called Deep Learning and the Future of AI, where it was pointed out that a machine’s ability to model the environment, predict possible futures and understand how the world works by observing it and acting in it, is still a very active research topic at the moment.

## 9 COMMERCIALY-AVAILABLE TECHNOLOGIES

Table 5 presents companies that offer a commercially available product; categorized by the major subjects of interest in this report. An attached spreadsheet provides additional information on each company, including what the tool uses deep learning to do, what type of deep learning is used and relevant notes, as available.

Table 5. Commercial Deep Learning Technologies

Category/Company	Category/Company	Category/Company
<b>Decision Support</b>	<b>Security</b>	<b>Medical Imaging/Diagnostics</b>
<i>Military:</i>	<i>Cyber:</i>	<a href="#">Atomwise</a>
<a href="#">Archarithms</a>	<a href="#">Cylance</a>	<a href="#">Avalon AI</a>
<a href="#">Massive Analytic</a>	<a href="#">Dark Trace</a>	<a href="#">Bay Labs</a>
<a href="#">Orbital Insight</a>	<a href="#">Deep Instinct</a>	<a href="#">Behold AI</a>
<a href="#">Stottler Henke Associates</a>	<a href="#">Demisto</a>	<a href="#">Butterfly Networks</a>
<a href="#">Video Inform</a>	<a href="#">Invincea</a>	<a href="#">Clarifai</a>
<i>Other:</i>	<a href="#">Ripjar</a> (cyber & audit)	<a href="#">Deep Genomics</a>
<a href="#">Digital Reasoning</a> (Various)	<a href="#">Securonix</a>	<a href="#">Enlitic</a>
<a href="#">Enlitic</a> (Medical)	<a href="#">Signal Sense</a>	<a href="#">Grail</a>
<a href="#">Gyana</a> (Urban Planning)	<a href="#">Vulnerability Insight</a>	<a href="#">Imagia</a>
<a href="#">Houzz</a> (Home Renovation)	<a href="#">Zimperium</a>	<a href="#">Lunit</a>
		<a href="#">Medy Match</a>

Category/Company	Category/Company	Category/Company
	<i>Physical:</i> <a href="#">Mentat</a> <a href="#">UMBO CV</a> <i>Other:</i> <a href="#">AppZen</a> (expense fraud)	<a href="#">Crixlabs</a> <a href="#">Recursion Pharmaceuticals</a> <a href="#">Zebra Medical Vision</a>
<b>Surveillance</b> <a href="#">Camio</a> <a href="#">Canary</a> <a href="#">DeepGlint</a> <a href="#">Herta Security</a> <a href="#">HyperVerge</a> <a href="#">Indicio</a> <a href="#">Orbital Insight</a> <a href="#">Salesforce Einstein (fka MetaMind)</a> <a href="#">ViSense</a>	<b>Emotion/Sentiment Recognition</b> <a href="#">Affectiva</a> <a href="#">Cogito</a> <a href="#">Ditto Labs</a> <a href="#">IBM Watson</a> <a href="#">Indico Data</a> <a href="#">Lexalytics</a>	<b>Big Data Analytics</b> <a href="#">Alphabyte Solutions</a> <a href="#">Arimo</a> <a href="#">Groundhog Technologies</a> <a href="#">IBM Watson</a> <a href="#">Pervazive</a> <a href="#">PointGrab</a>
<b>Event Detection</b> <a href="#">Ripjar</a> <a href="#">Tractable</a> <a href="#">UMBO CV</a>	<b>Recommender Systems</b> <a href="#">IBM Watson</a> <a href="#">Skymind.ai</a> <a href="#">Yusp (fka Gravity R&amp;D)</a>	<b>Risk Analysis</b> <a href="#">Intraspexion</a> (Legal) <a href="#">Legal Robot</a> (Legal) <a href="#">Resilinc</a> (Supply Chain) <a href="#">Vulnerability Insight</a> (Cybersec.)
<b>Forecasting</b> <a href="#">AlpacaAlgo</a> <a href="#">Shopper Trak</a>	<b>Data Fusion</b> <a href="#">GiantGrey</a> (SCADA) <a href="#">Leidos</a> (GEOINT)	<b>Other</b> <a href="#">deepsense.io</a> (deep learning project tracking)

## 10 CONCLUSIONS

The objective of this study was to detect and categorize the international R&D domains in the field of deep learning for human decision support, as well as to identify the commercially available products on the market. Results of this project can be used to identify domains that could present an area of interest for which DRDC may wish to develop expertise and to support the selection of future deep dive projects.

The current state of deep learning research, which has seen a significant jump in the last two years alone, is still largely focused on developing, modifying, adapting and testing various deep learning algorithms for different tasks and reporting its ability to surpass previous state-of-the-art alternatives. As Chris Nicholson, founder and CEO of the deep learning company Skymind, said in a recent interview “...most people are focused on the problems of making it work, applying it to existing problems, or inventing some incremental improvement”.<sup>254</sup> The improvements that are being achieved, however, are impressive. For example, Google’s algorithm, which learned to recognize cats in internet images, does so at a rate that is approximately 70% better than the previous state-of-the-art.

This report looked at two datasets related to deep learning; the broader set, which included 8,565 publications, and a human decision support subset containing 2,395 documents, which included client-selected topics from the full set. Thirteen different deep learning architectures were found in the dataset; however, each has a growing multitude of algorithmic variations. Convolutional neural networks (CNN), which is most frequently used for image related tasks, and deep neural networks

(DNN), which is most frequently used for speech related tasks, were the most frequently discussed in the literature both in the full and sub datasets.

Emerging trends were identified in the full dataset and described in more detail for the human decision making subset. Research momentum is picking up most notably in *Visualization, Surveillance, Event detection, Computer aided instruction, Recommender systems* and *Risk analysis*. Momentum has also been seen in *Decision making/support, Data fusion, Uncertainty analysis* and the *Military*. There was very limited evidence in the dataset that quantum computing is expected to have a major impact on deep learning. That being said, additional web searching and database scoping did reveal that there is a small amount of research in the area. As such, further investigation is recommended, in order to make a fuller assessment of the potential impact of quantum computing on deep learning.

The majority of the top affiliations are organizations from China, followed by the USA and Singapore. Three major corporations are amongst the top six affiliations, including Microsoft, Google and IBM. Similarly, three RTO affiliations are seen in the top 20, including the Chinese Academy of Sciences, which is the top affiliation in the full dataset, the Agency for Science, Technology and Research (A-STAR), Singapore and the Centre National de la Recherche Scientifique (CNRS), France. Canada, which is ranked 5<sup>th</sup> in the world in deep learning research, has expertise in *Imaging analysis & processing, Learning systems* and *Speech recognition*. One Canadian affiliation, the University of Toronto, also appears in the top 20 affiliations list. Analysis of the human decision support subset reveals that a number of the Canadian affiliations also have some capacity in *Forecasting/Prediction, Emotion/sentiment recognition, Human activity recognition, Searching, Diagnostics (medical), Recommender systems* and *Uncertainty analysis*.

A number of Canadian affiliations, including Concordia University, Université de Montreal (U de M), University of Toronto (U of T), University of British Columbia, York University and University of Guelph, are collaborating on the international scene. York University, UBC, U of T and U de M have collaborative relationships with IBM, Microsoft, Google or Facebook.

The human decision making subset is comprised of applications and tasks selected from the full set, suggesting that those were the primary applications areas in the human decision making subset. However, upon closer examination, additional applications were found in the subset, including, but not limited to, *Cybersecurity, Gaming, Simulation, and Virtual/Augmented reality*. This suggests that these domains also have decision making components that deep learning may be addressing. Additional analysis would be required for more details.

The current focus of deep learning research, i.e. significantly improving the algorithms within a short period of time, coupled with the fact that much of the research is being led by some major corporations, made identifying future applications and limitations challenging. There is minimal information in the literature and few market reports which discuss these issues in terms of deep learning for human decision support.

Sources that did report on future goals were more focused on the larger AI question regarding the potential for AI to surpass human intelligence and what consequences that might have for society. Frost & Sullivan projects that deep learning will impact aerospace and defence, amongst other sectors, by informing future impact scenarios in support of key decision makers. LeCun, the director of AI at Facebook, recently explained that building deep learning applications that integrate reasoning, have a

good architecture for short term memory, and can follow good principles for unsupervised learning are still far off on the horizon. Furthermore, he pointed out that a machine's ability to model the environment, predict possible futures and understand how the world works by observing and acting in it, is still a very active research topic at the moment. In recent interviews with some of the top names in deep learning, the most frequently mentioned area of research for the future is improvements in deep unsupervised learning. Gaining a fuller understanding of the specific challenges and keeping abreast of developments in this area is highly recommended. Similarly, it is highly recommended to explore IBM's "Cognitive Solutions", (outside the scope of this study), which IBM presents as having the potential to meet a decision support market that is projected to grow to roughly \$2 trillion USD by 2025.

In terms of deep learning limitations in general, and again, not specific to human decision support, the three main challenges are:

- Architectures and output models that are, in many ways, seen as a black box,
- Deep learning's inability to reason or apply learning to another context,
- The need for massive amounts of data.

In terms of deep learning using big data for decision support, the main outstanding issues relate to generalizability of training and learning from one domain to another.

Deep learning has the potential to dramatically improve the accuracy of data analytics for human decision support. However, modern deep learning research is still at a very early, albeit fast moving, stage of development. As such, this report provides a snapshot of the current state of research with some projections of future directions. It is highly recommended that this topic be followed closely for new developments as they are expected to occur at a fast pace. Furthermore, it is also recommended to include a subject specialist in future analyses of the domain.

## 11 REFERENCES

1. Najafabadi MM, Villanustre F, Khoshgoftaar TM, Seliya N, Wald R, Muharemagic E. Deep learning applications and challenges in big data analytics. *Journal of Big Data*. 2015;2(1):1.
2. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 05/28/print 2015;521(7553):436-444.
3. Benuwa BB, Zhan YZ, Ghansah B, Wornyo DK, Banaseka Kataka F. A Review of Deep Machine Learning. In: proceedings from International Journal of Engineering Research in Africa 2016.
4. Deng L. A tutorial survey of architectures, algorithms, and applications for deep learning. *APSIPA Transactions on Signal and Information Processing*. 2014;3:e2.
5. Marr B. A Short History of Deep Learning -- Everyone Should Read 2016; Available at: <http://www.forbes.com/sites/bernardmarr/2016/03/22/a-short-history-of-deep-learning-everyone-should-read/2/#63dcdf6e310f>. Accessed January 6, 2017.
6. Hernandez D. Meet the man Google hired to make AI a reality 2014; Available at: <https://www.wired.com/2014/01/geoffrey-hinton-deep-learning/>. Accessed January 6, 2017.
7. Nielsen MA. Neural Networks and Deep Learning: Determination Press; 2015: <http://neuralnetworksanddeeplearning.com/index.html>.
8. Karpathy A. CS231n Convolutional Neural Networks for Visual Recognition Available at: <http://cs231n.github.io/convolutional-networks/>. Accessed January 6, 2017.
9. Papernot N, McDaniel P, Jha S, Fredrikson M, Celik ZB, Swami A. The limitations of deep learning in adversarial settings. In: proceedings from 2016 IEEE European Symposium on Security and Privacy (EuroS&P) 2016.
10. Nielsen MA. Neural Networks and Deep Learning 2015; Available at: <http://neuralnetworksanddeeplearning.com/>. Accessed January 6, 2017.
11. Hinton GE, Osindero S, Teh Y-W. A fast learning algorithm for deep belief nets. *Neural Comput*. 2006;18.
12. Salakhutdinov R. Deep Learning: KDD Tutorial Available at: [http://www.cs.cmu.edu/~rsalakhu/talks/talk\\_KDD\\_part1\\_pdf.pdf](http://www.cs.cmu.edu/~rsalakhu/talks/talk_KDD_part1_pdf.pdf). Accessed January 6, 2017.
13. Smolensky P. *Information processing in dynamical systems: Foundations of harmony theory*: DTIC Document;1986.
14. Allahyar A. Bioinformatics: Collaboration of biology, math and computer science Available at: <http://prb.tudelft.nl/users/amin-allahyar>. Accessed January 6, 2017.
15. Hinton GE, Sejnowski TJ. Learning and relearning in Boltzmann machines. *Parallel distributed processing: Explorations in the microstructure of cognition*. 1986;1:282-317.
16. Salakhutdinov R, Hinton G. Deep Boltzmann Machines. AISTATS slides. 2009; Available at: <http://slideplayer.com/slide/6902409/>. Accessed January 6, 2017.
17. Salakhutdinov R, Hinton GE. Deep Boltzmann Machines. In: proceedings from AISTATS2009.
18. Roberts E. Neural Networks: Feed-Forward networks Available at: <http://cs.stanford.edu/people/eroberts/courses/soco/projects/neural-networks/Architecture/feedforward.html>. Accessed January 6, 2017.
19. Leverington D. A basic introduction to feedforward backpropagation neural networks 2009; Available at: [http://www.webpages.ttu.edu/dleverin/neural\\_network/neural\\_networks.html](http://www.webpages.ttu.edu/dleverin/neural_network/neural_networks.html). Accessed January 6, 2017.
20. Werbos P. Beyond regression: New tools for prediction and analysis in the behavioral sciences. 1974.
21. McClelland JL, Rumelhart DE, Hinton GE. The appeal of parallel distributed processing. *MIT Press, Cambridge MA*. 1986:3-44.

22. Deep Learning. DeepLearning 0.1 Documentation: Mulilayer Perceptron 2016; Available at: <http://deeplearning.net/tutorial/mlp.html>. Accessed January 6, 2017.
23. Honkela A. Nonlinear Switching State-Space Models 2001; Available at: <http://users.ics.aalto.fi/ahonkela/dippa/node41.html>. Accessed January 6, 2017.
24. Junior C, Moreira L, Rezende JLPd, Batista ALF, Mendonça ARd, Lacerda WS. Use of artificial neural networks for prognosis of charcoal prices in Minas Gerais. *Cerne*. 2013;19(2):281-288.
25. Bullinaria J. Recurrent Neural Networks. Neural Computation: Lecture 12 2015; Available at: <http://www.cs.bham.ac.uk/~jxb/INC/l12.pdf>. Accessed January 6, 2017.
26. Chiang Y-M, Chang L-C, Chang F-J. Comparison of static-feedforward and dynamic-feedback neural networks for rainfall-runoff modeling. *Journal of hydrology*. 2004;290(3):297-311.
27. Hochreiter S, Schmidhuber J. Long short-term memory. *Neural computation*. 1997;9(8):1735-1780.
28. Schmidhuber J. Why use recurrent networks at all? And why use a particular Deep Learning recurrent network called Long Short-Term Memory or LSTM? 2015; Available at: <http://people.idsia.ch/~juergen/rnn.html>. Accessed January 6, 2017.
29. Sutton C, McCallum A. An Introduction to Conditional Random Fields. *Foundations and Trends in Machine Learning*. 2011;4(5):267-373.
30. Yu D, Deng L, Wang S. Learning in the deep-structured conditional random fields. In: proceedings from NIPS Workshop 2009.
31. Do T, Arti T. Neural conditional random fields. In: proceedings from International Conference on Artificial Intelligence and Statistics 2010.
32. Yang X, Liu J. Deep belief network based CRF for spoken language understanding 2014.
33. Samek W, Binder A, Montavon G, Lapuschkin S, Muller K. Evaluating the Visualization of What a Deep Neural Network Has Learned. *IEEE Transactions on Neural Networks and Learning Systems*. 2016.
34. Khorrami P, Paine T, Huang T. Do deep neural networks learn facial action units when doing expression recognition? In: Proceedings of the IEEE International Conference on Computer Vision Workshops 2015.
35. Grün F, Rupprecht C, Navab N, Tombari F. A Taxonomy and Library for Visualizing Learned Features in Convolutional Neural Networks. *arXiv preprint arXiv:1606.07757*. 2016.
36. He K, Wang Y, Hopcroft J. A Powerful Generative Model Using Random Weights for the Deep Image Representation. In: proceedings from Advances In Neural Information Processing Systems 2016.
37. Xia Z, Zhu C, Wang Z, Guo Q, Liu Y. Every Filter Extracts A Specific Texture In Convolutional Neural Networks. *arXiv preprint arXiv:1608.04170*. 2016.
38. Yosinski J, Clune J, Nguyen A, Fuchs T, Lipson H. Understanding neural networks through deep visualization. *arXiv preprint arXiv:1506.06579*. 2015.
39. J. L, R. S, Wang. B, Qi. Y. Deep GDashboard: Visualizing and Understanding Genomic Sequences Using Deep Neural Networks. *arXiv* 2016;1608.03644v4.
40. Sim KC. Sensitivity-Characterised Activity Neurogram (SCAN) for Visualising and Understanding the Inner Workings of Deep Neural Network. *IEICE TRANSACTIONS on Information and Systems*. 2016;99(10):2423-2430.
41. Lee SH, Chan CS, Wilkin P, Remagnino P. Deep-plant: Plant identification with convolutional neural networks. In *Image Processing (ICIP), 2015 IEEE International Conference on* (pp. 452-456).

42. Mahendran A, Vedaldi A. Salient deconvolutional networks. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol 9910 LNCS: Springer Verlag; 2016:120-135.
43. Lo HZ, Cohen JP, Ding W. Prediction gradients for feature extraction and analysis from convolutional neural networks. In: Proceedings of the *Automatic Face and Gesture Recognition (FG), 2015 11th IEEE International Conference and Workshops on* (Vol. 1, pp. 1-6). IEEE.
44. Aghdam H, Heravi E, Puig D. Analyzing the stability of convolutional neural networks against image degradation. In: Proceedings of the 11th International Conference on Computer Vision Theory and Applications 2016.
45. Sinha Y. DFD: Data Flow Diagrams 2012; Available at: <http://www.slideshare.net/sinhamayur/data-flow-diagrams-11753554>. Accessed January 10, 2016.
46. Mahendran A, Vedaldi A. Visualizing Deep Convolutional Neural Networks Using Natural Pre-images. *International Journal of Computer Vision*. 2016;120(3):233-255.
47. Wang Y, Yao H, Zhao S. Auto-encoder based dimensionality reduction. *Neurocomputing*. 2016;184:232-242.
48. Aghdam HH, Heravi EJ, Puig D. Recognizing traffic signs using a practical deep neural network. In *Robot 2015: Second Iberian Robotics Conference* (Vol 417, pp. 399-410). Springer International Publishing.
49. Gandhi A, Poonia A. Image visualization as per the user appeal through deep learning. In *Proceedings of the ACM Symposium on Women in Research 2016* (pp. 81-86).
50. Wang Y, Dou Y, Lü Q, Li B, Li T. DLPF: A parallel deep learning programming framework based on heterogeneous architecture. *Jisuanji Yanjiu yu Fazhan/Computer Research and Development*. 2016;53(6):1202-1210.
51. Rossi B. How deep learning is advancing video analytics. 2016; Available at: <http://www.information-age.com/how-deep-learning-advancing-video-analytics-123461700/>. Accessed December 15, 2016.
52. Dashwood J. From recording to reacting: Neural networks are changing notions of surveillance 2016; Available at: <https://techcrunch.com/2016/04/14/from-recording-to-reacting-neural-networks-are-changing-notions-of-surveillance/>. Accessed December 15, 2016.
53. Du Y, Ai H, Lao S. A two-stage approach for bag detection in pedestrian images. In *Asian Conference on Computer Vision* (Vol 9006, pp. 507-521). Springer International Publishing.
54. Najva N, Bijoy KE. SIFT and Tensor Based Object Detection and Classification in Videos Using Deep Neural Networks. *Procedia Computer Science*. 2016;93:351-358.
55. Hayat M, Bennamoun M, An S. Deep reconstruction models for image set classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2015;37(4):713-727.
56. Peng X, Hoffman J, Yu SX, Saenko K. Fine-to-coarse Knowledge Transfer For Low-Res Image Classification. *arXiv preprint arXiv:1605.06695*. 2016.
57. Hou J, Wu X, Yu F, Jia Y. Multimedia event detection via deep spatial-temporal neural networks. In: proceedings from 2016 IEEE International Conference on Multimedia and Expo (ICME). 2016.
58. Luo S, Yang H, Wang C, Che X, Meinel C. Real-time action recognition in surveillance videos using ConvNets. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol 9949 LNCS: Springer Verlag; 2016:529-537.
59. Guo H, Wang J, Lu H. Multiple deep features learning for object retrieval in surveillance videos. *IET Computer Vision*. 2016;10(4):268-272.

60. Jin CB, Li S, Do TD, Kim H. Real-time human action recognition using CNN over temporal images for static video surveillance cameras. In *Pacific Rim Conference on Multimedia* (Vol 9315, pp. 330-339). Springer International Publishing. 2015.
61. Sarkar S. *Hierarchical symbolic perception in dynamic data driven application systems*: THE Pennsylvania State University; 2015.
62. Robinson JP, Fu Y. Pre-trained D-CNN models for detecting complex events in unconstrained videos. In: proceedings from SPIE Commercial+ Scientific Sensing and Imaging, 2016.
63. Minnehan B, Savakis A. Efficient pedestrian detection from aerial vehicles with object proposals and convolutional neural networks. In: proceedings from SPIE Defense+ Security, 2016.
64. Huo Y, Miao Z, Liu Z. A traffic object classification pipeline based on CNNs. In: proceedings from Chinese Automation Congress (CAC), 2015.
65. Li D, Chen X, Huang K. Multi-attribute learning for pedestrian attribute recognition in surveillance scenarios. In *Pattern Recognition (ACPR), 2015 3rd IAPR Asian Conference on* (pp. 111-115). IEEE.
66. Perlin HA, Lopes HS. Extracting human attributes using a convolutional neural network approach. *Pattern Recognition Letters*. 2015;68:250-259.
67. van de Wolfshaar J, Karaaba MF, Wiering MA. Deep Convolutional Neural Networks and Support Vector Machines for Gender Recognition. In: proceedings from Computational Intelligence, 2015 IEEE Symposium Series on, 2015.
68. John V, Mita S, Liu Z, Qi B. Pedestrian detection in thermal images using adaptive fuzzy C-means clustering and convolutional neural networks. In *Machine Vision Applications (MVA), 2015 14th IAPR International Conference on* (pp. 246-249). IEEE.
69. Liu H, Feng J, Qi M, Jiang J, Yan S. End-to-end comparative attention networks for person re-identification. *arXiv preprint arXiv:1606.04404*. 2016.
70. Varior RR, Haloi M, Wang G. Gated siamese convolutional neural network architecture for human re-identification. In: proceedings from European Conference on Computer Vision 2016.
71. Iodice S, Petrosino A, Ullah I. Strict pyramidal deep architectures for person re-identification. In *Advances in Neural Networks* (Vol 54, pp. 179-186). Springer International Publishing.
72. Filkovic I, Kalafatic Z, Hrkac T. Deep metric learning for person Re-identification and De-identification. In *Information and Communication Technology, Electronics and Microelectronics (MIPRO), 2016 39th International Convention on* (pp. 1360-1364). IEEE.
73. Chi H, Hu YH. Face de-identification using facial identity preserving features. In *Signal and Information Processing (GlobalSIP), 2015 IEEE Global Conference on* (pp. 586-590). IEEE.
74. Zaheer MZ, Kim JY, Kim H-G, Na SY. A Preliminary Study on Deep-Learning Based Screaming Sound Detection. In: proceedings from IT Convergence and Security (ICITCS), 2015 5th International Conference on.
75. Nikfarjam A, Sarker A, O'Connor k, Ginn R, Gonzalez G. Pharmacovigilance from social media: Mining adverse drug reaction mentions using sequence labeling with word embedding cluster features. *Journal of the American Medical Informatics Association*. 2015;22(3):671-681.
76. Ding C, Tao D. Trunk-Branch Ensemble Convolutional Neural Networks for Video-based Face Recognition. *arXiv preprint arXiv:1607.05427*. 2016.
77. Sarfraz MS, Stiefelhagen R. Deep Perceptual Mapping for Cross-Modal Face Recognition. *International Journal of Computer Vision*. 2016:1-13.
78. Zhou Y, Liu L, Shao L, Mellor M. DAVE: A unified framework for fast vehicle detection and annotation. In *European Conference on Computer Vision* (Vol 9906, pp. 278-293). Springer International Publishing.

79. Naimi A, Kessentini Y, Hammami M. Multi-nation and multi-norm license plates detection in real traffic surveillance environment using deep learning. In *International Conference on Neural Information Processing* (Vol. 9948, pp. 462-469). Springer International Publishing.
80. Liu X, Liu W, Ma H, Fu H. Large-scale vehicle re-identification in urban surveillance videos. In *Multimedia and Expo (ICME), 2016 IEEE International Conference on* (pp. 1-6). IEEE.
81. Zhou S, Shen W, Zeng D, Fang M, Wei Y, Zhang Z. Spatial-temporal convolutional neural networks for anomaly detection and localization in crowded scenes. *Signal Processing: Image Communication*. 2016;47:358-368.
82. Zhang C, Li H, Wang X, Yang X. Cross-scene crowd counting via deep convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 833-841).
83. Frost & Sullivan. *Advances in Surveillance Technology - Homeland Security TechVision Opportunity Engine*. D842-TV. 2016.
84. Developer N. Share your science: The future of facial recognition for video surveillance 2016; Available at: <https://news.developer.nvidia.com/share-your-science-the-future-of-facial-recognition-for-video-surveillance/>. Accessed December 15, 2016.
85. Ashraf K, Elizalde B, Iandola F, et al. Audio-based multimedia event detection with DNNs and sparse sampling. In *Proceedings of the 5th ACM on International Conference on Multimedia Retrieval*, pp. 611-614. ACM, 2015.
86. Revathi AR, Kumar D. An efficient system for anomaly detection using deep learning classifier. *Signal, Image and Video Processing*. 2016:1-9.
87. Al-Jarrah OY, Alhussein O, Yoo PD, Muhaidat S, Taha K, Kim K. Data Randomization and Cluster-Based Partitioning for Botnet Intrusion Detection. *IEEE Transactions on Cybernetics*. 2016;46(8):1796-1806.
88. Cakir E, Heittola T, Huttunen H, Virtanen T. Polyphonic sound event detection using multi label deep neural networks. In *Neural Networks (IJCNN), 2015 International Joint Conference on* (pp. 1-7). IEEE.
89. Jaccard N, Rogers TW, Morton EJ, Griffin LD. Automated detection of smuggled high-risk security threats using Deep Learning. *arXiv preprint arXiv:1609.02805*. 2016.
90. Hershey S, Chaudhuri S, Ellis DP, et al. CNN Architectures for Large-Scale Audio Classification. *arXiv preprint arXiv:1609.09430*. 2016.
91. Lam V, Phan S, Le DD, Duong DA, Satoh S. Evaluation of multiple features for violent scenes detection. *Multimedia Tools and Applications*. 2016:1-25.
92. McLoughlin I, Zhang H, Xie Z, Song Y, Xiao W. Robust sound event classification using deep neural networks. *IEEE Transactions on Audio, Speech and Language Processing*. 2015;23(3):540-552.
93. Espi M, Fujimoto M, Nakatani T. Acoustic event detection in speech overlapping scenarios based on high-resolution spectral input and deep learning. *IEICE Transactions on Information and Systems*. 2015;E98D(10):1799-1807.
94. Wang Y, Neves L, Metze F. Audio-based multimedia event detection using deep recurrent neural networks 2016.
95. Laffitte P, Sodayer D, Tatkeu C, Girin L. Deep neural networks for automatic detection of screams and shouted speech in subway trains. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on* (pp. 6460-6464). IEEE.
96. Espi M, Fujimoto M, Kinoshita K, Nakatani T. Exploiting spectro-temporal locality in deep learning based acoustic event detection. *EURASIP Journal on Audio, Speech, and Music Processing*, 2015(1), 26.

97. Zhang H, McLoughlin I, Song Y. Robust sound event recognition using convolutional neural networks. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on* (pp. 559-563). IEEE.
98. Dong B, Wang X. Comparison deep learning method to traditional methods using for network intrusion detection. In: proceedings from Communication Software and Networks (ICCSN), 2016 8th IEEE International Conference on 2016.
99. Gao N, Gao L, He YY. Deep belief nets model oriented to intrusion detection system. *Xi Tong Gong Cheng Yu Dian Zi Ji Shu/Systems Engineering and Electronics*. 2016;38(9):2201-2207.
100. Mohammadi S, Perina A, Kiani H, Murino V. Angry crowds: Detecting violent events in videos. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol 9911 LNCS: Springer Verlag; 2016:3-18.
101. Yuan Y, Jia K. A Distributed Anomaly Detection Method of Operation Energy Consumption Using Smart Meter Data. In *Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP), 2015 International Conference on* (pp. 310-313). IEEE
102. Chauhan S, Vig L. Anomaly detection in ECG time signals via deep long short-term memory networks. In *Data Science and Advanced Analytics (DSAA), 2015. 36678 2015. IEEE International Conference on* (pp. 1-7). IEEE.
103. Liu Y, Racah E, Correa J, et al. Application of Deep Convolutional Neural Networks for Detecting Extreme Weather in Climate Datasets. *arXiv preprint arXiv:1605.01156*. 2016.
104. Murphy RR. Computer vision image processing for finding missing persons. In proceedings from XPONENTIAL 2016 - An AUVSI Experience. 2016.
105. Besaw LE, Stimac PJ. Deep convolutional neural networks for classifying GPR B-scans. In: proceedings from SPIE Defense+ Security 2015.
106. Gan C, Wang N, Yang Y, Yeung D-Y, Hauptmann AG. Devnet: A deep event network for multimedia event detection and evidence recounting. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2015.
107. Hou J, Wu X, Yu F, Jia Y. Multimedia event detection via deep spatial-temporal neural networks. In *Multimedia and Expo (ICME), 2016 IEEE International Conference on* (pp. 1-6). IEEE.
108. Yu L, Sun X, Huang Z. Robust spatial-temporal deep model for multimedia event detection. *Neurocomputing*. 2016;213:48-53.
109. Akram SU, Kannala J, Eklund L, Heikkilä J. Cell segmentation proposal network for microscopy image analysis. In *International Workshop on Large-Scale Annotation of Biomedical Data and Expert Label Synthesis (Vol 10008, pp. 21-29)*. Springer International Publishing.
110. Sun W, Zheng B, Qian W. Computer aided lung cancer diagnosis with deep learning algorithms. In *SPIE Medical Imaging* (pp. 97850Z-97850Z). International Society for Optics and Photonics. 2016.
111. Qiu Y, Yan S, Tan M, Cheng S, Liu H, Zheng B. Computer-aided classification of mammographic masses using the deep learning technology: A preliminary study. In *SPIE Medical Imaging* (pp. 978520-978520). International Society for Optics and Photonics. 2016.
112. Benou A, Veksler R, Friedman A, Riklin Raviv T. De-noising of contrast-enhanced MRI sequences by an ensemble of expert deep neural networks. In *International Workshop on Large-Scale Annotation of Biomedical Data and Expert Label Synthesis (Vol 10008, pp. 95-110)*. Springer International Publishing. 2016.
113. Yoo Y, Tang LW, Brosch T, et al. Deep learning of brain lesion patterns for predicting future disease activity in patients with early symptoms of multiple sclerosis. In *International Workshop on Large-Scale Annotation of Biomedical Data and Expert Label Synthesis (Vol 10008, pp. 86-94)*. Springer International Publishing. 2016.

114. Jiao Z, Gao X, Wang Y, Li J. A deep feature based framework for breast masses classification. *Neurocomputing*. 2016;197:221-231.
115. Kisilev P, Sason E, Barkan E, Hashoul S. Medical image description using multi-task-loss CNN. In *International Workshop on Large-Scale Annotation of Biomedical Data and Expert Label Synthesis* (Vol 10008, pp. 121-129). 2016.
116. Suk HI, Shen D. Deep ensemble sparse regression network for alzheimer's disease diagnosis. In *International Workshop on Machine Learning in Medical Imaging* (Vol. 10019, pp. 113-121). Springer International Publishing. 2016.
117. Li W, Siniscalchi SM, Chen NF, Lee C-H. Improving non-native mispronunciation detection and enriching diagnostic feedback with DNN-based speech attribute modeling. In: proceedings from 2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)2016.
118. Elaraby MS, Abdallah M, Abdou S, Rashwan M. A Deep Neural Networks (DNN) Based Models for a Computer Aided Pronunciation Learning System. In: proceedings from International Conference on Speech and Computer2016.
119. Smith A, Aksit O, Min W, Wiebe E, Mott BW, Lester JC. Integrating real-time drawing and writing diagnostic models: An evidence-centered design framework for multimodal science assessment. In *International Conference on Intelligent Tutoring Systems* (Vol 9684pp. 165-175). Springer International Publishing. 2016.
120. Hu B, Terrazas BV. Building a Mental Health Knowledge Model to Facilitate Decision Support. In: proceedings from Pacific Rim Knowledge Acquisition Workshop.2016.
121. Bansal T, Belanger D, McCallum A. Ask the GRU: Multi-task learning for deep text recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems* (pp. 107-114). ACM. 2016.
122. Chen L, Wang JF, Wang N, Li X. DNN based mobile online video recommendation. *Jisuanji Xuebao/Chinese Journal of Computers*. 2016;39(8):1626-1638.
123. Zhang W, Du T, Wang J. Deep learning over Multi-Field categorical Data - A case study on user response prediction. In *European Conference on Information Retrieval*, (Vol 9626, pp. 45-57). Springer International Publishing, 2016.
124. Li X, Tian J, Xu M, Ning Y, Cai L. DBLSTM-based multi-scale fusion for dynamic emotion prediction in music. In *Proceedings - IEEE International Conference on Multimedia and Expo*. 2016.
125. Guo G. Resolving data sparsity and cold start in recommender systems. In *Proceedings International Conference on User Modeling, Adaptation, and Personalization*. 2012.
126. Wei J, He J, Chen K, Zhou Y, Tang Z. Collaborative Filtering and Deep Learning Based Hybrid Recommendation for Cold Start Problem. In *Proceedings from Dependable, Autonomic and Secure Computing, 14th Intl Conf on Pervasive Intelligence and Computing, 2nd Intl Conf on Big Data Intelligence and Computing and Cyber Science and Technology Congress (DASC/PiCom/DataCom/CyberSciTech)*, 2016 IEEE.
127. Zhang F, Yuan NJ, Lian D, Xie X, Ma WY. Collaborative knowledge base embedding for recommender systems. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 353-362). 2016.
128. Janvier E, Couronne T, Grozavu N. Features learning and transformation based on deep autoencoders. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol 9949 LNCS: Springer Verlag; 2016:111-118.
129. Kim D, Park C, Oh J, Lee S, Yu H. Convolutional matrix factorization for document context-aware recommendation. In *Proceedings of the 10th ACM Conference on Recommender Systems*, pp. 233-240. ACM, 2016.

130. Zhang W, Du T, Wang J. Deep Learning over Multi-field Categorical Data. In: proceedings from European Conference on Information Retrieval. In Proceedings *European Conference on Information Retrieval*, pp. 45-57. Springer International Publishing, 2016.
131. Musto C, Greco C, Suglia A, Semeraro G. Ask Me Any Rating: A content-based recommender system based on recurrent neural networks. In Proceedings CEUR Workshop (Vol. 1653). 2016.
132. Wu S, Ren W, Yu C, Chen G, Zhang D, Zhu J. Personal recommendation using deep recurrent neural networks in NetEase. In *Data Engineering (ICDE), 2016 IEEE 32nd International Conference on*, pp. 1218-1229. IEEE, 2016.
133. Unger M, Shapira B, Rokach L, Bar A. Deep auto-encoding for context-aware inference of preferred items' categories. In Proceedings CEUR Workshop (Vol. 1688).2016.
134. Unger M, Bar A, Shapira B, Rokach L. Towards latent context-aware recommendation systems. *Knowledge-Based Systems*.2016;104:165-178.
135. Ahn Y, Lee H, Jeon H, Ha S, Lee SG. Quote recommendation for dialogs and writings. *CBRecSys 2016*, 39. 2016.
136. Gadepally VN, Hancock BJ, Greenfield KB, Campbell JP, Campbell WM, Reuther AI. Recommender Systems for the Department of Defense and Intelligence Community. *Lincoln Laboratory Journal*. 2016;22(1).
137. Moy G. Machine-Learning & Recommender Systems for C2 of Autonomous Vehicles 2016; Available at: [http://www.dst.defence.gov.au/sites/default/files/events/documents/ISTAS-2016\\_Machine\\_Learning\\_Recommender\\_Systems\\_Glenn-Moy.pdf](http://www.dst.defence.gov.au/sites/default/files/events/documents/ISTAS-2016_Machine_Learning_Recommender_Systems_Glenn-Moy.pdf). Accessed January 12, 2017.
138. Sanchez-Riera J, Hua KL, Hsiao YS, Lim T, Hidayati SC, Cheng WH. A comparative study of data fusion for RGB-D based visual recognition. *Pattern Recognition Letters*. 2016;73:1-6.
139. Zhang P, Ma X, Zhang W, et al. Multimodal fusion for sensor data using stacked autoencoders. In Proceedings *Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), 2015 IEEE Tenth International Conference on* (pp. 1-2). IEEE. 2015.
140. Tao X, Kong D, Wei Y, Wang Y. A big network traffic data fusion approach based on fisher and deep auto-encoder. *Information (Switzerland)*. 2016;7(2).
141. Gandhi A, Sharma A, Biswas A, Deshmukh O. GeThR-Net: A Generalized Temporally Hybrid Recurrent Neural Network for Multimodal Information Fusion. In: proceedings from European Conference on Computer Vision (pp. 883-899). 2016.
142. Li Z, Gan Y, Liang X, Yu Y, Cheng H, Lin L. LSTM-CF: Unifying context modeling and fusion with LSTMs for RGB-D scene labeling. In Proceedings from *European Conference on Computer Vision* (Vol 9906, pp. 541-557). Springer International Publishing. 2016.
143. Giering M, Venugopalan V, Reddy K. Multi-modal sensor registration for vehicle perception via deep neural networks. In *High Performance Extreme Computing Conference (HPEC), 2015 IEEE* (pp. 1-6).
144. Molchanov P, Gupta S, Kim K, Pulli K. Multi-sensor system for driver's hand-gesture recognition. In *Automatic Face and Gesture Recognition (FG), 2015 11th IEEE International Conference and Workshops on*, vol. 1, pp. 1-8. IEEE, 2015.
145. Jain A, Singh A, Koppula HS, Soh S, Saxena A. Recurrent Neural Networks for driver activity anticipation via sensory-fusion architecture. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on* (pp. 3118-3125). IEEE, 2016.
146. Liu S, Liu S, Cai W, et al. Multimodal Neuroimaging Feature Learning for Multiclass Diagnosis of Alzheimer's Disease. *IEEE Transactions on Biomedical Engineering*. 2015;62(4):1132-1140.
147. Yin XX, Zhang Y, Cao J, Wu JL, Hadjiloucas S. Exploring the complementarity of THz pulse imaging and DCE-MRIs: Toward a unified multi-channel classification and a deep learning framework. *Computer Methods and Programs in Biomedicine*. 2016;137:87-114.

148. Xu Y, Li Y, Liu M, et al. Gland instance segmentation by deep multichannel side supervision. In: proceedings from International Conference on Medical Image Computing and Computer-Assisted Intervention, 2016.
149. Tucker P. The Pentagon Needs Its Own Google For All Its Data, Says Eric Schmidt 2017; Available at: <http://www.defenseone.com/technology/2017/01/pentagon-needs-its-own-google-all-its-data-says-eric-schmidt/134456/?oref=d-channelriver>. Accessed January 19, 2017.
150. Defense Innovation Board. Fact Sheet on Recommendations for the Public Meeting on January 9, 2017. 2017; Available at: [https://www.defense.gov/Portals/1/Documents/pubs/DIB\\_Recommendations\\_Executive\\_Summary\\_170106.pdf](https://www.defense.gov/Portals/1/Documents/pubs/DIB_Recommendations_Executive_Summary_170106.pdf). Accessed January 19, 2017.
151. Gal Y. *Uncertainty in Deep learning*. U.K.: Machine Learning Group, Department of Engineering, University of Cambridge; 2016.
152. Grzeszick R, Sudholt S, Fink GA. Optimistic and Pessimistic Neural Networks for Scene and Object Recognition. *arXiv preprint arXiv:1609.07982*. 2016.
153. Li C, Stevens A, Chen C, Pu Y, Gan Z, Carin L. Learning weight uncertainty with stochastic gradient MCMC for shape classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5666-5675. 2016.
154. Zhang N, Ding S, Zhang J, Xue Y. Research on Point-wise Gated Deep Networks. *Applied Soft Computing Journal. (Preprint)* 2016.
155. Srinivasan S, Wang D. A supervised learning approach to uncertainty decoding for robust speech recognition. In: proceedings from Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on 2006.
156. Tachioka Y, Watanabe S. Uncertainty training and decoding methods of deep neural networks based on stochastic representation of enhanced features. In *INTERSPEECH*, pp. 3541-3545. 2015.
157. Abdelaziz AH, Watanabe S, Hershey JR, Vincent E, Kolossa D. Uncertainty propagation through deep neural networks. In *Interspeech* 2015.
158. Astudillo RF, Correia J, Trancoso I. Integration of DNN based Speech Enhancement and ASR. In *INTERSPEECH* (pp. 3576-3580). 2015.
159. Huemmer C, Maas R, Schwarz A, Astudillo RF, Kellermann W. Uncertainty decoding for DNN-HMM hybrid systems based on numerical sampling. In *INTERSPEECH* (pp. 3556-3560). 2015.
160. Mallidi SH, Ogawa T, Hermansky H. Uncertainty estimation of DNN classifiers. In Proceedings from *In Automatic Speech Recognition and Understanding (ASRU), 2015 IEEE Workshop on* (pp. 283-288). IEEE. 2015.
161. Qiu Y, Wang Y, Yan S, et al. An initial investigation on developing a new method to predict short-term breast cancer risk based on deep learning technology. In *SPIE Medical Imaging* (pp. 978521-978521). International Society for Optics and Photonics. 2016.
162. Liu S, Liu S, Cai W, Pujol S, Kikinis R, Feng D. Early diagnosis of Alzheimer's disease with deep learning. In proceedings from *Biomedical Imaging (ISBI), 2014 IEEE 11th International Symposium on* (pp. 1015-1018). IEEE. 2014.
163. Futoma J, Morris J, Lucas J. A comparison of models for predicting early hospital readmissions. *Journal of Biomedical Informatics*. 2015;56:229-238.
164. Miotto R, Li L, Dudley JT. Deep learning to predict patient future diseases from the electronic health records. In *European Conference on Information Retrieval (Vol 9626, pp. 768-774)*. Springer International Publishing. 2016.

165. Yu L, Yang Z, Tang L. A novel multistage deep belief network based extreme learning machine ensemble learning paradigm for credit risk assessment. *Flexible Services and Manufacturing Journal*. 2015.
166. Ha VS, Nguyen HN. Credit scoring with a feature selection approach based deep learning. In *MATEC Web of Conferences* (Vol. 54). EDP Sciences. 2016.
167. Heaton JB, Polson NG, Witte JH. Deep learning for finance: Deep portfolios. *Applied Stochastic Models in Business and Industry*. 2016.
168. Miyazaki H, Kuwata K, Ohira W, et al. Development of an automated system for building detection from high-resolution satellite images. In *Proceedings from Earth Observation and Remote Sensing Applications (EORSA), 2016 4th International Workshop on* (pp. 245-249). IEEE. 2016.
169. Sivakorn S, Polakis I, Keromytis AD. I am Robot: (Deep) learning to break semantic image CAPTCHAs. In *proceedings from Security and Privacy (EuroS&P), 2016 IEEE European Symposium on* (pp. 388-403). IEEE. 2016.
170. Chen J, Tang X. The challenges and feasibility of societal risk classification based on deep learning of representations. In *Systems, Man, and Cybernetics (SMC), 2015 IEEE International Conference on*, pp. 569-574. IEEE, 2015.
171. Wan K, Alagar V. Risk and vulnerability analysis of critical infrastructure. In *proceedings from International Conference on Intelligent Computing* (Vol 9773, pp. 54-66). Springer International Publishing. 2016.
172. Hu C, Ju R, Shen Y, Zhou P, Li Q. Clinical decision support for Alzheimer's disease based on deep learning and brain network. In *Communications (ICC), 2016 IEEE International Conference on* (pp. 1-6). IEEE. 2016.
173. Stier N, Vincent N, Liebeskind D, Scalzo F. Deep learning of tissue fate features in acute ischemic stroke. In *proceedings from Bioinformatics and Biomedicine (BIBM), 2015 IEEE International Conference on* (pp. 1316-1321). IEEE 2015.
174. Seguí S, Drozdal M, Pascual G, et al. Generic feature learning for wireless capsule endoscopy analysis. *Computers in Biology and Medicine*. 2016;79:163-172.
175. Hu B, Terrazas BV. Building a mental health knowledge model to facilitate decision support. In *proceedings from Pacific Rim Knowledge Acquisition Workshop* (Vol 9806, pp. 198-212). Springer International Publishing. 2016.
176. Hu Z, Hu J, Ding W, Zheng X. Review Sentiment Analysis Based on Deep Learning. In *proceedings from e-Business Engineering (ICEBE), 2015 IEEE 12th International Conference on*, pp. 87-94. IEEE, 2015.
177. Zhao X, Wang C, Yang Z, Zhang Y, Yuan X. Online news emotion prediction with bidirectional LSTM. In *proceedings from International Conference on Web-Age Information Management* (Vol 9659, pp. 238-250). Springer International Publishing. 2016.
178. Li L, Ren W, Qin B, Liu T. Learning document representation for deceptive opinion spam detection. In *proceedings from Chinese Computational Linguistics and Natural Language Processing Based on Naturally Annotated Big Data* (Vol 9427, pp. 393-404). Springer International Publishing. 2015.
179. Zhang H, Sanin C, Szczerbicki E. When Neural Networks Meet Decisional DNA: A Promising New Perspective for Knowledge Representation and Sharing. *Cybernetics and Systems*. 2016;47(1-2):140-148.
180. Leng J, Jiang P. A deep learning approach for relationship extraction from interaction context in social manufacturing paradigm. *Knowledge-Based Systems*. 2016;100(188-199).

181. Lin KY, Tsai JJP. A deep learning-based customer forecasting tool. In proceedings from *Multimedia Big Data (BigMM), 2016 IEEE Second International Conference on* (pp. 198-205). IEEE. 2016.
182. Zhu F, Liu Q, Zhong S, Yang Y. Human-level moving object recognition from traffic video. *Computer Science and Information Systems*. 2015;12(2):787-799.
183. Jiang P, Chen C, Liu X. Time series prediction for evolutions of complex systems: A deep learning approach. In proceedings from *Control and Robotics Engineering (ICCRE), 2016 IEEE International Conference on* (pp. 1-6). IEEE. 2016.
184. Alam F, Mehmood R, Katib I, Albeshri A. Analysis of Eight Data Mining Algorithms for Smarter Internet of Things (IoT). *Procedia Computer Science* 2016;98:437-442.
185. News Medical. UCSF collaborates with Intel to deploy and validate deep learning analytics to improve healthcare 2017; Available at: <http://www.news-medical.net/news/20170119/UCSF-collaborates-with-Intel-to-deploy-and-validate-deep-learning-analytics-to-improve-healthcare.aspx>. Accessed January 20, 2017.
186. Transparency Market Research. Healthcare Natural Language Processing (NLP) Market Outlook 2024, Advancement in Machine Language 2017; Available at: <http://www.transparencymarketresearch.com/pressrelease/healthcare-natural-language-processing-market.htm>. Accessed January 20, 2017.
187. Forbes. First FDA Approval For Clinical Cloud-Based Deep Learning In Healthcare 2017; Available at: <http://www.forbes.com/sites/bernardmarr/2017/01/20/first-fda-approval-for-clinical-cloud-based-deep-learning-in-healthcare/#126c7f1e46e6>. Accessed January 20, 2017.
188. Webwire. Intel Experts Share Insights on Tech Trends for 2017 and Beyond 2016; Available at: <http://www.webwire.com/ViewPressRel.asp?ald=207082>. Accessed January 20, 2017.
189. Henke N, Bughin J, Chui M, et al. The Age of Analytics: Competing in a Data-Driven World 2016; Available at: <http://www.mckinsey.com/~media/mckinsey/business%20functions/mckinsey%20analytics/our%20insights/the%20age%20of%20analytics%20competing%20in%20a%20data%20driven%20world/mgi-the-age-of-analytics-executive-summary.ashx>. Accessed January 20, 2017.
190. Gold M. *Explorer: Artificial Intelligence*: Strategic Business Insights;2016.
191. Konkel F. The CIA Says It Can Predict Social Unrest as Early as 3 to 5 Days Out 2016; Available at: [http://www.defenseone.com/technology/2016/10/cia-says-it-can-predict-social-unrest-early-3-5-days-out/132121/?oref=search\\_deep%20learning](http://www.defenseone.com/technology/2016/10/cia-says-it-can-predict-social-unrest-early-3-5-days-out/132121/?oref=search_deep%20learning). Accessed January 20, 2017.
192. Becker J, Havens TC, Pinar A, Schulz TJ. Deep belief networks for false alarm rejection in forward-looking ground-penetrating radar. In proceedings from *SPIE Defense+ Security* (pp. 94540W-94540W). International Society for Optics and Photonics. 2015.
193. Besaw LE, Stimac PJ. Deep convolutional neural networks for classifying GPR B-Scans. In proceedings from *SPIE Defense+ Security* (pp. 945413-945413). International Society for Optics and Photonics. 2015.
194. Wu H, Zhang H, Zhang J, Xu F. Typical Target Detection in Satellite Images Based on Convolutional Neural Networks. In proceedings from *Systems, Man, and Cybernetics (SMC), 2015 IEEE International Conference on*, pp. 2956-2961. IEEE, 2015.
195. Wei J, Liu CH, Zhu Z, Mendoza-Schrock O, Vongsy K. Classification of uncooperative vehicles with sparse laser Doppler vibrometry measurements. In *SPIE Defense+ Security* (pp. 94640Y-94640Y). International Society for Optics and Photonics. 2015.
196. Qu Y, Jiang L, Guo X. Moving vehicle detection with convolutional networks in UAV videos. In proceedings from *Control, Automation and Robotics (ICCAR), 2016 2nd International Conference on* (pp. 225-229). IEEE. 2016.

197. Cao L, Jiang Q, Cheng M, Wang C. Robust vehicle detection by combining deep features with exemplar classification. *Neurocomputing*. 2016;215:225-231.
198. Zhang F, Du B, Zhang L, Xu M. Weakly Supervised Learning Based on Coupled Convolutional Neural Networks for Aircraft Detection. *IEEE Transactions on Geoscience and Remote Sensing*. 2016;54(9):5553-5563.
199. Weisgerber M. The Supercomputer That Can Spot Mobile Missiles. 2016; Available at: <http://ai.business/2016/06/21/artificial-intelligence-in-defence-and-security-industry/>. Accessed January 13, 2017.
200. Zhu F, Hu X, Wu L, He X, Guo S. Inspiration for battlefield situation cognition from AI military programs launched by DARPA of USA and development of AI technology. *Communications in Computer and Information Science*. 2016;644:566-577.
201. Contract Research and Development. R&D Business Opportunity Abstracts. 2016; Available at: <http://www.contractrandd.com/?s=deep+learning>. Accessed January 12, 2017.
202. IARPA. Creation of Operationally Realistic 3D Environment (CORE3D). 2016; Available at: <https://www.iarpa.gov/index.php/research-programs/core3d?highlight=WyJkZWVwliwibGVhcm5pbmciLCJkZWVwIGxlYXJuaW5nIl0=>. Accessed January 13, 2017.
203. IARPA. Deep Intermodal Video Analytics (DIVA). 2016; Available at: <https://www.iarpa.gov/index.php/research-programs/diva>. Accessed January 13, 2017.
204. DARPA. Explainable Artificial Intelligence (XAI). 2016; Available at: <http://www.darpa.mil/attachments/DARPA-BAA-16-53.pdf>. Accessed January 13, 2017.
205. U.S. Defense Advanced Research Projects Agency. Research, Development, Test and Evaluation (RDT&E): Volume 1. Washington, DC: DARPA; 2016: <http://www.darpa.mil/attachments/DARPAFY17PresidentsBudgetRequest.pdf>
206. Frost & Sullivan. *Big Data in Defense and Security: From the User's Perspective*. 9831-16.2015.
207. Tucker P. New Microchip Could Increase Military Intelligence Powers Exponentially. 2016; Available at: <http://www.defenseone.com/technology/2016/02/new-microchip-could-increase-military-intelligence-powers-exponentially/125715/>. Accessed January 13, 2017.
208. Eidam E. Pentagon Eyes Deep Machine Learning in Fight Against ISIS. 2016; Available at: <http://www.govtech.com/federal/Pentagon-Eyes-Deep-Machine-Learning-in-Fight-Against-ISIS.html>. Accessed January 13, 2017.
209. Wiebe N, Kapoor A, Svore KM. Quantum deep learning. *Quantum Information and Computation*. 2016;16(7-8):541-587.
210. Adachi SH, Henderson MP. Application of quantum annealing to training of deep neural networks. *arXiv preprint arXiv:1510.06356*. 2015.
211. D-Wave. Careers: Machine learning Scientist 2016; Available at: <http://www.dwavesys.com/careers/machine-learning-scientist>. Accessed January 20, 2017.
212. Novet J. Google Researchers: Quantum Computers Aren't Perfect for Deep Learning. 2015; Available at: <http://venturebeat.com/2015/11/11/google-researcher-quantum-computers-arent-perfect-for-deep-learning>. Accessed January 20, 2017.
213. University of Montreal. Yoshua Bengio. Available at: [http://www.iro.umontreal.ca/~bengio/yoshua\\_en/index.html](http://www.iro.umontreal.ca/~bengio/yoshua_en/index.html). Accessed January 5, 2017.
214. University of Montreal. MILA: Montreal Institute for Learning Algorithms. 2016; Available at: <https://mila.umontreal.ca/en/mila/>. Accessed January 5, 2017.
215. Michalski V, Memisevic R, Konda K. Modeling deep temporal dependencies with recurrent "Grammar Cells". In proceedings from *Advances in neural information processing systems* (pp. 1925-1933).2014.

216. Denil M, Shakibi B, Dinh L, Ranzato M, De Freitas N. Predicting parameters in deep learning. In proceedings from *Advances in Neural Information Processing Systems* (pp. 2148-2156). 2013.
217. Goodfellow IJ, Mirza M, Courville A, Bengio Y. Multi-prediction deep Boltzmann machines. In proceedings from *Advances in Neural Information Processing Systems* (pp. 548-556). 2013.
218. Konda K, Memisevic R. Learning visual odometry with a convolutional network. 2015. Available at: <https://www.iro.umontreal.ca/~memisevr/pubs/VISAPP2015.pdf>. Accessed January 24, 2017.
219. Kahou SE, Pal C, Bouthillier X, et al. Combining modality specific deep neural networks for emotion recognition in video. In *Proceedings of the 15th ACM on International conference on multimodal interaction* (pp. 543-550). ACM. 2013.
220. Kahou SE, Bouthillier X, Lamblin P, et al. EmoNets: Multimodal deep learning approaches for emotion recognition in video. *Journal on Multimodal User Interfaces*. 2016;10(2):99-111.
221. Kahou SE, Michalski V, Konda K, Memisevic R, Pal C. Recurrent neural networks for emotion recognition in video. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction* (pp. 467-474). ACM. 2015.
222. Rifai S, Bengio Y, Courville A, Vincent P, Mirza M. Disentangling factors of variation for facial expression recognition. In proceedings from *European Conference on Computer Vision*, (Vol 7577, pp. 808-822). Springer Berlin Heidelberg, 2012.
223. Glorot X, Bordes A, Bengio Y. *Domain adaptation for large-scale sentiment classification: A deep learning approach*. In *Proceedings of the 28th international conference on machine learning (ICML-11)* (pp. 513-520). 2011.
224. Martinez HP, Bengio Y, Yannakakis G. Learning deep physiological models of affect. *IEEE Computational Intelligence Magazine*. 2013;8(2):20-33.
225. Shen Y, He X, Gao J, Deng L, Mesnil G. A latent semantic model with convolutional-pooling structure for information retrieval. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management* (pp. 101-110). ACM. 2014.
226. Shen Y, He X, Gao J, Deng L, Mesnil G. Learning semantic representations using convolutional neural networks for web search. In *Proceedings of the 23rd International Conference on World Wide Web* (pp. 373-374). ACM. 2014.
227. Bergstra J, Bengio Y. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*. 2012;13:281-305.
228. Ding X, Liu T, Duan J, Nie JY. Mining user consumption intention from social media using domain adaptive convolutional neural network. In proceedings of the 29<sup>th</sup> AAAI Conference on Artificial Intelligence. (pp.2389-2395). AAAI. 2015.
229. University of Toronto. Geoffrey E. Hinton Available at: <http://www.cs.toronto.edu/~hinton/>. Accessed January 5, 2017.
230. Zheng Y, Zemel RS, Zhang YJ, Larochelle H. A Neural Autoregressive Approach to Attention-based Recognition. *International Journal of Computer Vision*. 2015;113(1):67-79.
231. Salakhutdinov R, Tenenbaum JB, Torralba A. Learning to learn with compound HD models. In proceedings from *Advances in Neural Information Processing Systems* (pp. 2061-2069). 2011.
232. Salakhutdinov R, Tenenbaum JB, Torralba A. Learning with hierarchical-deep models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2013;35(8):1958-1971.
233. Ranzato M, Mnih V, Susskind JM, Hinton GE. Modeling natural images using gated MRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2013;35(9):2206-2222.
234. Jaitly N, Vanhoucke V, Hinton G. Autoregressive product of multi-frame predictions can improve the accuracy of hybrid models. In *INTERSPEECH* (pp. 1905-1909). 2014.

235. Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*. 2014;15:1929-1958.
236. Alipanahi B, Delong A, Weirauch MT, Frey BJ. Predicting the sequence specificities of DNA- and RNA-binding proteins by deep learning. *Nature Biotechnology*. 2015;33(8):831-838.
237. Bai M, Luo W, Kundu K, Urtasun R. Exploiting semantic information and deep matching for optical flow. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. 2016; 9910:154-170.
238. Peretroukhin V, Clement L, Kelly J. Reducing Drift in Visual Odometry by Inferring Sun Direction using a Bayesian Convolutional Neural Network. *arXiv preprint arXiv:1609.05993*. 2016.
239. University of British Columbia. Computer Science Research Groups Available at: <https://www.cs.ubc.ca/cs-research/lci/research-groups>. Accessed January 5, 2017.
240. Palangi H, Ward R, Deng L. Using deep stacking network to improve structured compressed sensing with Multiple Measurement Vectors. 2013. Available at: <https://www.microsoft.com/en-us/research/publication/using-deep-stacking-network-to-improve-structured-compressive-sensing-with-multiple-measurement-vectors/> Accessed: January 24, 2017.
241. Azizi S, Imani F, Zhuang B, et al. Ultrasound-based detection of prostate cancer using automatic feature selection with deep belief networks. In proceedings from *International Conference on Medical Image Computing and Computer-Assisted Intervention (Vol 9350, pp. 70-77)*. Springer International Publishing. 2015.
242. Suzani A, Seitel A, Liu Y, Fels S, Rohling RN, Abolmaesumi P. Fast Automatic Vertebrae Detection and Localization in Pathological CT Scans-A Deep Learning Approach. In proceedings from *International Conference on Medical Image Computing and Computer-Assisted Intervention*. (pp. 678-686). 2015.
243. Brosch T, Tam R. Efficient training of convolutional deep belief networks in the frequency domain for application to high-resolution 2D and 3D images. *Neural Computation*. 2015;27(1):211-227.
244. Allen K. How a Toronto professor's research revolutionized artificial intelligence. 2015; Available at: <https://www.thestar.com/news/world/2015/04/17/how-a-toronto-professors-research-revolutionized-artificial-intelligence.html>. Accessed January 5, 2017.
245. CIFAR. CIFAR names Geoffrey Hinton a Distinguished Fellow 2015; Available at: <https://www.cifar.ca/assets/cifar-names-geoffrey-hinton-a-distinguished-fellow/>. Accessed January 5, 2017.
246. Honari S, Yosinski J, Vincent P, Pal C. Recombinator networks: Learning coarse-to-fine feature aggregation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 5743-5752). 2016.
247. Kotzias D, Denil M, De Freitas N, Smyth P. From group to individual labels using deep features. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 597-606). ACM. 2015.
248. Neverova N, Wolf C, Paci G, Somnavilla G, Taylor GW, Nebout F. A multi-scale approach to gesture detection and recognition. In *Proceedings of the IEEE International Conference on Computer Vision Workshops* (pp. 484-491). 2013.
249. Neverova N, Wolf C, Taylor GW, Nebout F. Multi-scale deep learning for gesture detection and localization. In proceedings from *Workshop at the European Conference on Computer Vision (Vol 8925, pp. 474-490)*. Springer International Publishing. 2015.
250. Neverova N, Wolf C, Taylor G, Nebout F. ModDrop: Adaptive multi-modal gesture recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2016;38(8):1692-1706.

251. Neverova N, Wolf C, Taylor GW, Nebout F. Hand segmentation with structured convolutional learning. In: proceedings from Asian Conference on Computer Vision. (pp. 687-702). 2014.
252. Frost & Sullivan. *Deep Learning Interfaces: Future of Artificial Intelligence, D66B-TI2015*.
253. LeCun Y. Deep Learning and the Future of AI 2016; Available at: <https://indico.cern.ch/event/510372/attachments/1245509/1840815/lecun-20160324-cern.pdf>. Accessed January 24, 2017.
254. Motley J. We chat with deep learning company, Skymind, about the future of AI. 2016; Available at: <http://thenextweb.com/insider/2016/12/23/we-chat-with-deep-learning-company-skymind-about-the-future-of-ai/>. Accessed January 17, 2017.
255. Curtis S. What to Expect from Deep Learning in 2016 and Beyond 2016; Available at: <http://www.kdnuggets.com/2016/01/deep-learning-2016-beyond.html>. Accessed January 17, 2017.
256. Hemsoth N. The Next Wave of Deep Learning Applications 2016; Available at: <https://www.nextplatform.com/2016/09/14/next-wave-deep-learning-applications/>. Accessed January 17, 2017.
257. Zhang B. Big Geospatial Data + Deep Learning + High Performance Computing = Geospatial Intelligence 2016; Available at: <http://on-demand.gputechconf.com/gtc/2016/presentation/s6260-bingcai-zhang-big-geospatial-data-deep-learning-high-performance-computing.pdf>. Accessed January 17, 2017.
258. IBM. IBM Investor Briefing 2016; Available at: [https://www.ibm.com/investor/att/pdf/2016\\_Investor\\_Briefing\\_Cognitive\\_Solutions.pdf](https://www.ibm.com/investor/att/pdf/2016_Investor_Briefing_Cognitive_Solutions.pdf). Accessed January 17, 2017.
259. NESTA. Machines that Learn in the Wild 2015; Available at: [https://www.nesta.org.uk/sites/default/files/machines\\_that\\_learn\\_in\\_the\\_wild.pdf](https://www.nesta.org.uk/sites/default/files/machines_that_learn_in_the_wild.pdf). Accessed January 13, 2017.
260. Ambati S. Deep learning: A brief guide for practical problem solvers. 2015. Available at: <http://www.infoworld.com/article/3003315/big-data/deep-learning-a-brief-guide-for-practical-problem-solvers.html>. Accessed January 24, 2017.
261. Vincent J. These are three of the biggest problems facing today's AI 2016; Available at: <http://www.theverge.com/2016/10/10/13224930/ai-deep-learning-limitations-drawbacks>. Accessed January 13, 2017.
262. Melendez S. Deep Learning: Teaching Computers to Predict the Future 2016; Available at: <http://www.nbcnews.com/mach/technology/deep-learning-predicts-future-n690851>. Accessed January 17, 2017.
263. Frost & Sullivan. *Big Data in Defence and Security: From the User's Perspective. 9831-16*. 2015.
264. Knight W. New Tool Lets AI Learn to Do Almost Anything on a Computer 2016; Available at: <https://www.technologyreview.com/s/603011/new-tool-lets-ai-learn-to-do-almost-anything-on-a-computer/>.
265. Strachan-Olsen D. *In Deep Learning, Quantity Matters*. Strategic Business Insights, Soc899. 2016.
266. Najafabadi MM, Villanustre F, Khoshgoftaar TM, Seliya N, Wald R, Muharemagic E. Deep learning applications and challenges in big data analytics. *Journal of Big Data*. 2015;2(1):1-21.
267. Deng L. Deep Learning for AI from Machine Perception to Machine Cognition 2016; Available at: [http://www.icassp2016.org/SP16\\_PlenaryDeng\\_Slides.pdf](http://www.icassp2016.org/SP16_PlenaryDeng_Slides.pdf). Accessed January 17, 2017.

## 12 APPENDIX A: ATTACHMENTS

The following files have been provided as attachments to the recipients of this report.

**Table 6. Attachments**

Filename	Description
1. Deep Learning Visualizations	Tableau workbook
2. Deep Learning Companies	Excel

## 13 APPENDIX B: METHODOLOGY

### 13.1 Searches

To address the key questions in this project, a broad search on next generation computing was conducted in several databases including *Scopus*, *National Technical Information Service (NTIS)* and *Inspec* between 2011 and 2016. The terms listed in Table 7 were used in multiple combinations with various Boolean operators depending on the database that was searched. An additional search was conducted in *Defense Technical Information Center (DTIC)* to support the description of research trends in the *Military* subject group but the results were not part of the dataset that was analyzed in all other sections of the report.

**Table 7. Search Terms**

Generic	Known Architectures
"deep neural networks"	"deep belief networks"
"deep learning"	"convolutional neural networks"
"deep structured learning"	"deep boltzmann machines"
"hierarchical learning"	"deep auto-encoders"
"deep machine learning"	"deep stacking networks"
"deep memory networks"	"compound hierarchical deep models"
deep reinforcement learning	"multilayer kernel machines"
	"deep Q-networks"
	deep q-learning
	deep convolutional network
	convolutional deep network
	deep autoencoder
	convnet
	lstm OR "long short-term memory" AND deep
	restricted boltzmann machine AND deep
	recurrent neural networks AND deep

## 13.2 Analysis

All references were downloaded into VantagePoint software for analysis. VantagePoint enables the creation of various groupings, statistical analyses, matrices, graphs, and cross-correlations to analyze the data and profile the activities of the major players.

Keywords, classification codes, index keywords, key phrase identifiers, subject headings, author keywords and words and phrases extracted by natural language processing of titles were merged together to facilitate subject analysis (hereafter referred to as keywords). Keywords were then organized into subject groups through a variety of steps.

Different analytical tools were used to generate graphs based on statistical operations performed in VantagePoint. TouchGraph software was used for cluster analysis and visualization of the subject groups while Tableau software was used to generate bubble graphs.

## 13.3 R&D Momentum

The R&D Momentum indicator is designed to identify rapidly rising subjects with relatively few publications. The challenge of identifying such subjects lies with the publication volume as a confounding factor, for their rapid growth and evolution is dwarfed by the high volume of established subjects. Specifically, the notion of "emerging" consists not only of a sharply rising trend line but also of a small footprint in the domain of interest. A *relatively* small footprint is the reason emerging subjects are often overlooked until their disruptive impacts become obvious. In the Momentum indicator, the two parameters correspond to (1) growth rate which is the slope of a subject's trend line (right-left axis), and (2) volume which is the cumulated total number of publications (vertical axis).

Once growth rate and volume are separated, a two-dimensional coordinate can be used to plot a group of subjects. To do so, the two parameters have to be normalized with z-scores. The normalization process converts two sets of values in different units into the same measure by means of standard deviation, which also standardizes the variations for each of the two parameters. The four-quadrant visualization provides a structured view of the relative position of these subjects within the group.

The techniques involved in producing this indicator can be applied to both literature and patent data. However, emergence is a concept more relevant to scientific breakthroughs than to patents, because the latter is closer to the application stage when the technology of concern has matured. The literature covers a broader range of topics, is more diverse, and explores pre-commercial problems by attempting different strategies. This sort of exploration usually precedes what eventually shows up in patents, whereas patents represent a relatively small subset of what is seen in the literature, i.e., a subset of what can be exploited commercially. Therefore, we may assume that emerging subjects are much more likely to appear in the literature. Once reaching the patenting stage, they have generally passed the exploration phase and have already demonstrated real or anticipated commercial attractiveness.

DOCUMENT CONTROL DATA		
(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)		
1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.)  NRC-Knowledge Management 1200 Montreal Rd., Ottawa, ON K1A 0R6	2. SECURITY CLASSIFICATION (Overall security classification of the document including special warning terms if applicable.)  UNCLASSIFIED	
3. TITLE Deep Learning for Human Decision Support		
4. AUTHORS Erica Wiseman, PhD		
5. DATE OF PUBLICATION  January 2017	6a. NO. OF PAGES  70p.	6b. NO. OF REFS (Total cited in document.)  267
7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)		
8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)  Defence R&D Canada Ottawa, Ontario Canada		
9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)	9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)	
10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.)	10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)	
11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)  Unlimited		
12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.)  Unlimited		