Center for Digital Strategies
Tuck at Dartmouth

Attivio: Machine Learning for Strategic Insights[1]

Stephen Baker, CEO of Attivio was having a great day. His company Attivio was just recognized by Forrester as one of the leaders in cognitive search, a relatively new term to describe enterprise search solutions that employ AI technologies such as natural language processing and machine learning to understand, organize, and query digital content from disparate internal and external sources of data and content.[2] Baker had recently received project inquiries from several Fortune 500 companies. The project that intrigued Baker the most was the opportunity to work with HSBC– one of the world's largest banking and financial services institutions. HSBC inquired if Attivio could solve its problems with identification of financial crime and fraud by its customers. The issue was of utmost importance to the bank in 2017 as it had just overcome recent sanctions due to failure to detect such incidents. The project was exciting but Baker wanted to be sure to continue the successful momentum of the company and wondered if the problems faced by HSBC were in the sweet spot for Attivio's methodologies and technologies. The company could not afford to slow its string of achievements or allocate resources that could be used on one of the many other project opportunities that the recent positive press was generating.

ARTIFICIAL INTELLIGENCE – BORN AT DARTMOUTH

Famed computer scientist John McCarthy first coined the term artificial intelligence in 1956 at a conference at Dartmouth College. At its core, AI refers to a machine that can replicate any level of human intelligence. While there are theoretically increasingly complex levels of AI, from programs that address a specific area (often called narrow intelligence - ANI) to self-thinking systems found in science fiction movies (e.g. Skynet from the Terminator movie series), most of the time when AI is used in the business, it is referring to narrow intelligence ANI systems. Multiple ANI systems can work together within a single piece of technology building what appears to be a more advanced form of AI, but at its core each individual system performs a specific task. For example, Google Translate is an application that takes spoken user input in one language and delivers a translation. The first ANI piece of the language translation app is the voice recognition software that uses natural language processing to determine what the user is saying. Next, a separate ANI system is used to find the most appropriate translation match in the desired output language that is ultimately fed back to the user. Self-driving cars are filled with perhaps hundreds of different ANIs working together to determine optimal routes, detect and avoid hazards on the road, set

---

[1] Professor Alva Taylor, Stephen Baker and Michael Hanley (T'19 and MBA Associate of the Center for Digital Strategies) prepared this Tuck School of Business case solely as the basis for class discussion. Cases are not intended to serve as endorsements, sources of primary data, or illustrations of effective or ineffective management.

[2] Source: Forrester Report: The Forrester Wave™: Cognitive Search and Knowledge Discovery Solutions

and maintain safe following distances, and all the other necessary actions required for a car to truly drive itself.

While AI refers to the system that executes the specific action at or above human capabilities, the real magic behind the scenes is machine learning (ML). ML is the basic practice of using algorithms to parse data, learn from it, and then make a determination or prediction. So rather than hand-coding software routines with a specific set of instructions to accomplish a particular task, the machine is "trained" using large amounts of data that give it the ability to learn how to perform the task.

Big Data: Unlocking Machine Learning's Power

While the concept of AI has been around since the mid-1950s, the technical capabilities have been out of reach until relatively recently. Enabling AI so that the system can perform as well as humans on specific tasks requires that it be trained on a large dataset. The term **big data** is used to describe any kind of data set that has a least one of four shared characteristics, called the four Vs:

- Extremely large *Volumes* of data
- The ability to move that data at a high *Velocity* of speed
- An ever-expanding *Variety* of data sources
- *Veracity* so that data sources represent the underlying truth

AI needs to be able to gather, store, manage and manipulate vast amounts of evolving data in a timely manner to provide the right insights. Amazon's Alexa constantly needs to adapt to stay relevant to the modern consumer's constantly changing preferences and behaviors. To do so, Alexa requires large datasets to learn from so that the AI can perform new skills and continually improve existing ones.

Machine Learning: How AI Knows What to Do

AI requires large amounts of data if it hopes to be accurate in executing tasks as well as a human. But how does an AI system interact with big data? Simply put, what does it mean when a machine learns? There are four major "learning" techniques used to train an AI system: *Supervised Learning, Unsupervised Learning, Reinforced Learning,* and *Deep Learning*.

- **Supervised Learning:** Supervised learning begins with an established set of preprocessed data and a certain understanding of how the data is classified. This dataset is used as the training dataset to build the correlations between classification variables through statistical regression. The idea behind supervised learning is that you provide the AI with enough known examples of something that it will be able to accurately classify a never-seen-before example of what it has been trained. One major problem of supervised learning is *overfitting* which occurs when

an AI model is too precisely tuned to the training dataset and ends up being inaccurate with unknown datasets that don't perfectly resemble the training data. Increasing the size of the training data set with examples that make the overall data set less homogenous can reduce overfitting.

- **Unsupervised Learning:** Unlike supervised learning, this technique is used when presented with a large set of data that is unlabeled. Unsupervised learning algorithms segment data into groups of examples (clusters) or groups of features by looking for patterns in the data. The algorithms then create the parameter values and classifications that were lacking in the data that in turn makes the data supervised. Unsupervised learning can help businesses understand large volumes of new, unlabeled data where classification by hand would prove to be extremely costly and time-consuming. For example, in healthcare, collecting huge amounts of unlabeled data about a specific disease can help practitioners gain insights into the patterns of symptoms and relate those to outcomes from patients.
- **Reinforced Learning:** Reinforced learning is a behavioral learning model in which algorithms receive feedback from the analysis of the data so that it is guided towards the best answer. Unlike other types of learning techniques, reinforced learning does not start with a sample data set. Instead, it learns through a trial-and-error process in which the successful decisions are reinforced in the algorithm. This technique is popular with robotics and game playing. Actions that result in failure, such as the loss of a game, are avoided in the next trial and the process is repeated until the AI has a successful sequence of actions to follow under given circumstances. AlphaGo, an AI developed by Alphabet Inc.'s Google DeepMind that plays the popular board game Go, was initially trained using supervised and unsupervised learning to mimic moves from expert players from historical games. To improve beyond this proficiency the AI was trained further by being set to play large numbers of games against itself, reinforcing moves that resulted in higher probabilities of winning.
- **Deep Learning:** Deep learning, also called complex neural networks, are designed to emulate how the human brain works so computers can be trained to deal with abstractions and problems that are poorly defined. Deep learning is a machine learning technique that uses hierarchical neural networks to learn from a combination of unsupervised and supervised algorithms. A neural network consists of three or more layers: an input layer, one or many hidden layers, and an output layer. Data is ingested through the input layer. Then the data is modified in the hidden layer and the output layers based on the weights applied to these layers from the layers lower down in the hierarchy. Using an iterative approach, a neural network continuously adjusts and makes inferences until a specific stopping point is reached. Neural networks are often used for image recognition and computer vision applications. For an example of how deep learning algorithms work see Exhibit 1,

which illustrates how a computer vision program can be used to identify a person based on an image of their face.

Training AI in any of the above ways requires significant amounts of training data whether it be structured, unstructured, or in the form of trial-and-error. Use of AI has increased in recent years thanks to the explosion of data storage capabilities through the cloud, increased processing power through GPU advancement, and data availability through the propagation of Internet of Things (IoT) devices that capture data as actions are performed. Some examples of already available AI products and services include:

- Personal assistants such as Amazon's Alexa are able to digest spoken language and perform actions based on the request.
- Spotify can classify a user based on their listening history and make recommendations of new music based on listening patterns of similar users.
- Has Facebook ever asked you, unprompted, if a picture contains you and it does? Thanks to face-detection AI trained on other tagged photos of you and others, Facebook is now able to make predictions about who might be in a photo.
- GE is using machine learning and AI to create systems that detect when mechanical malfunctions are likely to occur based on historical malfunction reports. Combined with real-time data from its products, GE is able to predict problems before they occur to improve maintenance practices.

COMPANY AND COMPETITOR BACKGROUND - ATTIVIO

Founded in 2007, and based in Boston, Attivio has grown to approximately 90 employees. The company provides customer insight and search solutions to primarily Fortune 500 clients using industry-leading natural language processing, machine learning, analytics, and knowledge graphing capabilities in an overall solution often described as cognitive search. Cognitive search promised significant advantages over traditional search, particularly when dealing with vast amounts of complex data.

Attivio uses AI to improve content curation through identification of documents, which leads to improved search relevancy. Instead of simply returning all documents that contain search keywords, Attivio is able to provide search results given the context of the query based on similar searches from the past. Important documents that flow through a company's system can be automatically identified and brought to the attention of users for immediate action. Imagine instead of having to sift through your inbox at the beginning of the day to determine what tasks you need to prioritize, you are shown right away which action items deserve your immediate attention and which can be tackled later.

Table 1. Cognitive Search vs. Traditional Enterprise Search

	Enterprise Search	Cognitive Search
↑ Demand for Answers	# ↑ $ ↑	# ↓ $ ↑
↑ Technology for Answers	# ↑ $ ↓	# ↓ $ ↑

Recently, Attivio's solutions were rated as one of the leaders in the cognitive search field, and one that has exhibited a strong implementation and execution strategy.

Table 2: Forrester Analysis of Cognitive Search Competitive Landscape

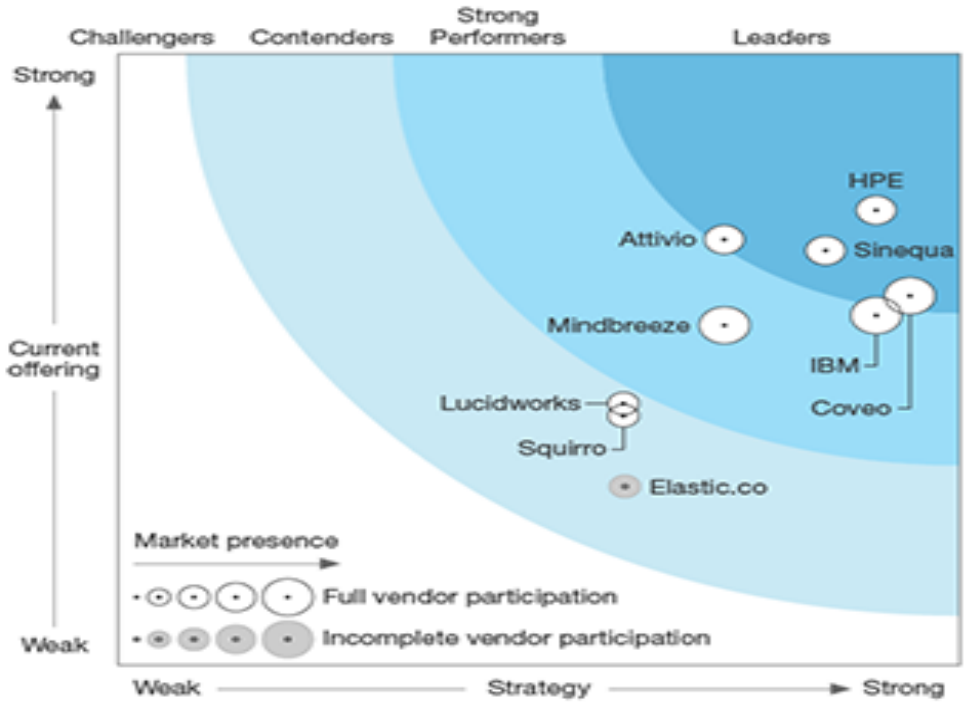

Attivio target markets are strategic search deployments that are inherently large, complex, and mission critical. These types of deployments are typically found in the Fortune 1000/Global 2000, thus the company targets these organizations and structures their marketing and sales plans to drive awareness and capture demand.

Within the target market, Attivio specifically targets the persona of the Enterprise Architect (and equivalent roles within the IT department). This persona can be described as "bilingual" in both technology and business with 82% of Enterprise Architects designing and building IT strategies to support business needs. To achieve this, the typical Enterprise Architect spends 20 – 25% of their time researching new technologies and vendors. Buy-in from this persona is critical for a successful engagement. In Attivio's view, the Enterprise Architect (EA) often has three primary needs:

Employee Productivity: The EA's company has employees that need to be able to answer complex questions, by providing immediate insight across all available information. EAs need fast and secure solutions that can scale as her or his company grows. Typical needs include Knowledge Management, Research Portal, Corporate Intranet, and Site Search.

Customer Support: EA's are often responsible for providing customer support and service, including call center operations. Solutions are needed to access and deliver comprehensive views of customer and product/service information regardless of communication channel. Typical applications include Contact Center Search, Self-Service Support, and Preventative Maintenance.

Risk Avoidance: EA's, in the role of the company's steward for information, also require systems that assure compliance and detection of activities that place the company's information at risk. Typical applications include Communication Surveillance and Customer Due Diligence. Key functional requirements include the ability to monitor and analyze hundreds-of-billions of communications and transactions against compliance policies and normalize/resolve entities with accuracy.

Competitor Approaches

Attivio typically faces competition from two general sources: commercial vendors of Cognitive Search/ML platforms (IBM, HP, Sinequa); and from open source technologies. IBM's Watson is an example of a commercial vendor that offers a platform that provides similar capabilities to Attivio. IBM draws on it well-established brand and deep industry expertise in competing for projects. In most cases, IBM is more expensive and complex in its implementation than Attivio, with most projects lasting more than a year and typical project costs greater than $10 million for a combination of software and hardware. [3]

Open source solutions from smaller platforms such as Lucene, Solr, and Elastic are another competitive approach to Attivio. These open solutions tend to be more basic, drawing on widely available techniques making them more suitable for basic information retrieval use cases, versus full-blown machine learning. All of the fundamental capabilities are developed, maintained, and advanced by a massive community of developers. One

[3] For additional information see: https://www.technologyreview.com/s/607965/a-reality-check-for-ibms-ai-ambitions/

advantage of these solution approaches is that they are much less costly than those of commercial vendors, as they don't require on-going licensing or maintenance fees[4].

THE HSBC PROJECT

Like most financial institutions, since 2012 HSBC was facing increased regulatory and legal demands to comply with anti-money laundering (AML) and countering funding of terrorism (CFT). During that time period, regulators raised their expectations on financial institutions to prevent and detect suspicious transactions. The cost of adhering to these expectations was growing rapidly. In 2013, JPMorgan added 4,000 employees to their compliance team and spent an additional $1 billion on controls, while other companies like Citicorp and UBS spent an estimated $2 billion and $1 billion respectively on new compliance spending. During that time, HSBC increased the personnel in their compliance department from 2,000 to 5,000 and continued this increase to over 7,000 by 2016. It is estimated that 30% of banking staff are required to meet the governance, risk, and compliance (GRC) needs of banks.[5] Also, during that time period, Bain & Co. reported that GRC accounted for 15 to 20% of the total bank operating costs for most major banks.[6]

The increased financial and personnel resources were used in an increased effort to identify and investigate suspicious transactions. As regulations have increased, financial institutions are not only required to identify questionable transactions but to also provide evidence that they have done appropriate levels of due diligence and research on both the transaction and any and all parties involved in the transactions. Typical investigation processes are manual with each flagged transaction requiring the development of a narrative containing over 100 facts about the parties involved. An individual case can take over an hour to review and document. The top 20 banks flag more than 150, 000 transactions per month.

Penalties for failing to adhere to the increased standards are quite severe. HSBC has been hit particularly hard by failures to prevent financial crime. In 2012 the company was fined almost $1.5B for allowing Mexican drug traffickers to make illegal deposits. Not only was the company subject to this significant fine, but the violation also resulted in higher levels of compliance requirements called Deferred Prosecution Agreement (DPA). Estimated total costs for this violation including oversight, additional insurance, and fines was $5 billion.[7] HSBC was not the only bank to face these fees and increased compliance oversight. The British Bank Barclays was fined over $350 million, and after the announcement of the fine its stock price fell 15% (-$4.5B) and its CEO was forced to resign. The stock price dropped not only due to the fines, but also in response to anticipated future cost, plus the reputation and trust lost by customers of and investors in the bank. The Standard

[4] For additional information see: https://www.cioinsight.com/it-strategy/infrastructure/slideshows/five-pros-and-five-cons-of-open-source-software)
[5] *Are Compliance Costs Breaking Banks?* Trulioo.com, August 25, 2015
[6] *You've Heard of Fintech, Get Ready for 'Regtech'*, BankThink, September 2016.
[7] For additional information see https://www.theguardian.com/business/2017/may/04/hsbc-first-quarter-pre-tax-profit-falls-19

Chartered Bank was fined $340 million for lax compliance efforts that made the financial system vulnerable to money laundering.[8]

Recently, HSBC faced another challenge as the lawyer charged with guiding the additional oversight of the company under the DPA announced that the company was once again being investigated for violating money laundering compliance laws. On the heels of this announcement, company shares dropped by 62% in 2016.

In partial response to the public announcement of the additional compliance reviews, HSBC's CEO Stuart Gulliver said the bank was unearthing more regulatory problems due to higher-quality internal policing, not new issues. He stated that additional violations were identified because of their diligence in finding bad actors amongst their 37 million customer base. He said it was "quite normal" to uncover such instances in a bank the size of HSBC which operates in 70 countries and has 240,000 staff.[9]

Methods of Money Laundering

There are three typical stages involved in money laundering: placement, layering and integration (see Exhibit 2 for an illustration of a typical money laundering flow).[10]

Placement –This is the movement of cash from its source. On occasion the source can be easily disguised or misrepresented. The cash is then placed into circulation through financial institutions, casinos, shops, currency converters, and other businesses. The process of placement can be carried out through many processes including:

1. *Currency Smuggling* – This is the physical illegal movement of currency and monetary instruments out of a country.
2. *Bank Complicity* – This is when a financial institution, such as banks, is owned or controlled by unscrupulous individuals who cooperate with drug dealers and other organized illegal groups.
3. *Currency Exchanges* – In a number of foreign exchange markets, currency movements can support the illegal movement of currency.
4. *Securities Brokers* – Brokers can facilitate the process of money laundering through structuring large deposits of cash in a way that disguises the original source of the funds.
5. *Blending of Funds* – Money from illicit activities is used to set up front companies. This enables the funds from illicit activities to be obscured in legal transactions.
6. *Asset Purchase* – The purchase of assets with cash with the major purpose to change the form of the proceeds from conspicuous bulk cash to some equally valuable but less conspicuous form.

[8] *HSBC to pay £1.2bn over Mexico scandal*, The Guardian, December 2012.

[9] *City watchdog investigating HSBC over potential financial crime, bank reveals*, The Guardian, Feb. 2017

[10] The source of much of the descriptions in this section are from *Money Laundering* in the UK, Ronojit Banerjee, and Guy Stessens, *Money Laundering*, (2000 Cambridge University Press, Cambridge)

Layering – The purpose of this stage is to make it more difficult to detect and uncover a laundering activity. It is meant to make the trailing of illegal proceeds difficult for law enforcement agencies. The known methods are:

1. *Cash converted into Monetary Instruments* – Once the placement is successful within the financial system by way of a bank or financial institution, the proceeds can then be converted into monetary instruments. This involves the use of banker's drafts and money orders.
2. *Material assets bought with cash then sold* – Assets that are bought through illicit funds can be resold locally or abroad and in such a case the assets become more difficult to trace and thus seize.

Integration – This is the movement of previously laundered money into the economy mainly through the banking system and thus such monies appear to be normal business earnings. Typical methods used are:

1. *Property Dealing* – The sale of property to integrate laundered money back into the economy, many criminal groups use shell companies to buy property with proceeds from the sale considered legitimate.
2. *Front Companies* – Front companies that are incorporated in countries with corporate secrecy laws, in which criminals lend themselves their own laundered proceeds in an apparently legitimate transaction.
3. *Foreign Bank Complicity* – Money laundering using known foreign banks represents a higher order of sophistication and presents a very difficult target for law enforcement. The willing assistance of the foreign banks is frequently protected against law enforcement scrutiny. This is not only through criminals but also by banking laws and regulations of other sovereign countries.
4. *False Import/Export Invoices* – The use of false invoices by import/export companies is an effective way of integrating illicit proceeds back into the economy. This involves the overvaluation of entry documents to justify funds later deposited in domestic banks and/or the value of funds received from exports.

Banks are incurring massive costs and employing hundreds of people to monitor transactions and investigate accounts, but they can only identify potential problems based off of pre-defined rules and patterns.

HSBC Reaches Out to Attivio

A key challenge that faced HSBC and any large bank remained: How to cost-effectively prevent financial crimes given the current regulatory environment? Getting in a position where you are being investigated by one of the regulatory bodies has to be avoided, otherwise there is significant equity, reputation, and operating cost risk. At the same time, risk operations are a cost-center, eat dramatically into the bottom line, and are characterized by manual, costly workflows. As a bank's business grows, so do the number

of transactions, communications, products that introduce risk, and the associated potential cost of financial crimes prevention.

HSBC was reaching the end of the 5-year deferred prosecution agreement that was instituted in 2012 and wanted to avoid further problems and issues that would result in further penalties and oversight. The bank's management looked to ML solutions to help combat the problems with their current practices. One issue with the traditional approach used by HSBC and other financial institutions for anti-money laundering (AML) compliance purposes is that it results in a high rate of false positives. As concerns have increased, valid transactions and valid customers are being investigated resulting in wasted resources. Another issue with HSBC's current AML operations is that violators and money launderers are constantly changing and inventing new ways to try to circumvent detection.

The Decision

Baker viewed the HSBC opportunity with both excitement and caution. It would be a huge win for the company to help the Bank solve the AML issues. However, a failed project would stall the impressive growth of the company, and sully its current reputation. His question was whether he could take on the project? How should it be structured? And how should he manage the expectations of HSBC management?

Exhibit 1:
Steps to Identify a Person via Deep Learning Algorithms[11]

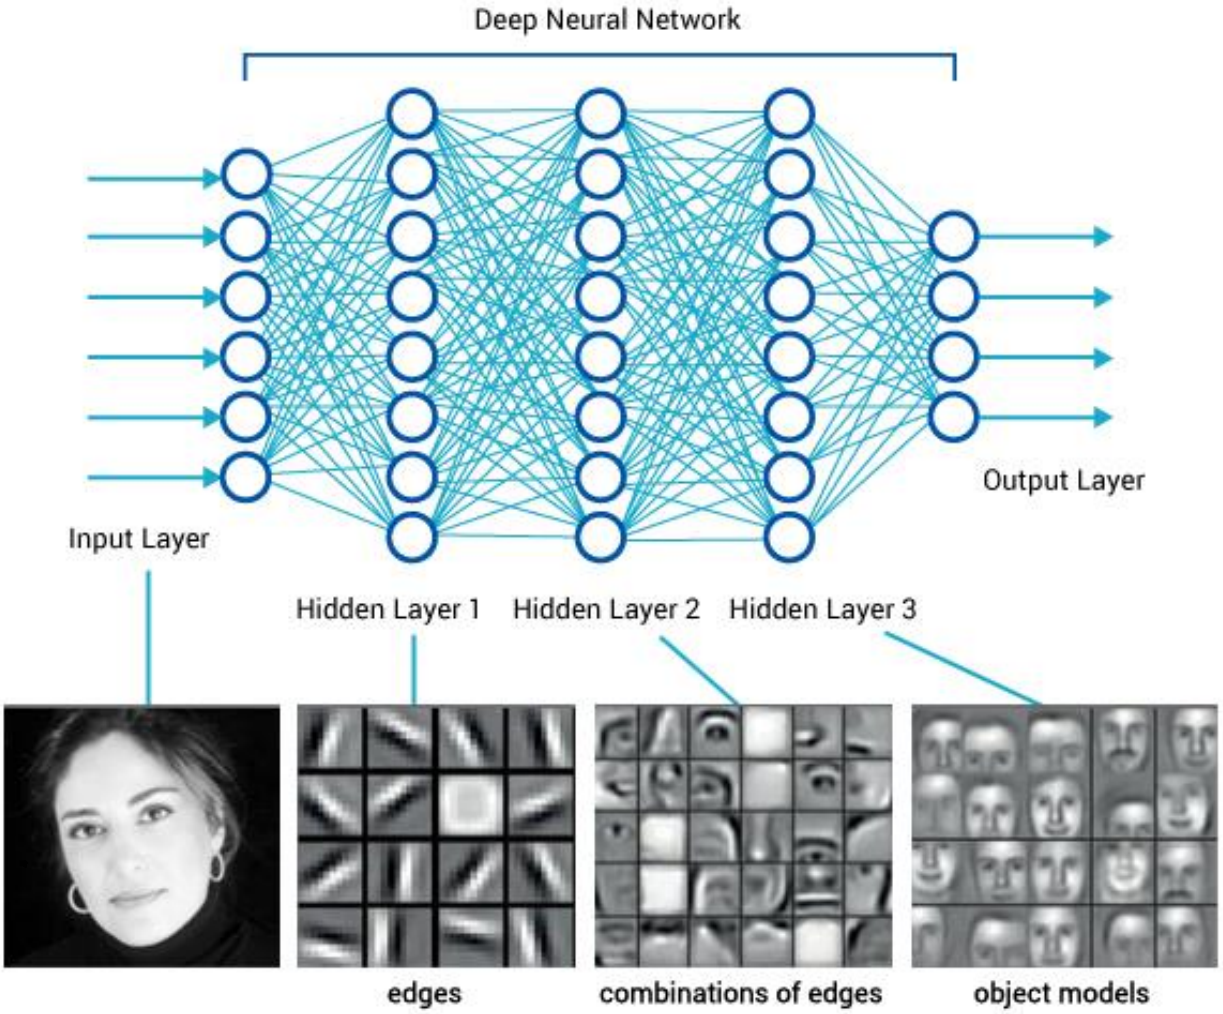

1. The input layer is created by breaking down the original image into smaller grid pieces which become the input nodes for the rest of the neural network.

2. The first hidden layer is made up of edge detection algorithms that have been trained to identify and classify specific edges. Each node in the layer represents a different edge detection algorithm.

[11] **Source: Deep Learning Smarts Up Your Smart Phone, December 8, 2015, R. Meyer, https://www.amax.com/blog/**

Exhibit 1 (continued)
Steps to Identify a Person via Deep Learning Algorithms

3. The second hidden layer takes input from the first hidden layer in the form of the detected edges and identifies combinations of edges that it has been trained to find.

4. The third hidden layer goes one step further and identifies object models from the combination of edges identified in the second hidden layer.

5. The final output layer could be images that are similar to the original input layer or even the person's name if the object models had been classified. This is the process Facebook uses to auto-tag photos of you that had been uploaded without being tagged. Facebook's algorithms have been trained to iterate through the above process and compare the object models to already tagged photos. Each output, as there can be multiple, has a probability of match associated with it that if above a set threshold will trigger the automatic tagging procedure.

Exhibit 2: An Example of Money Laundering System

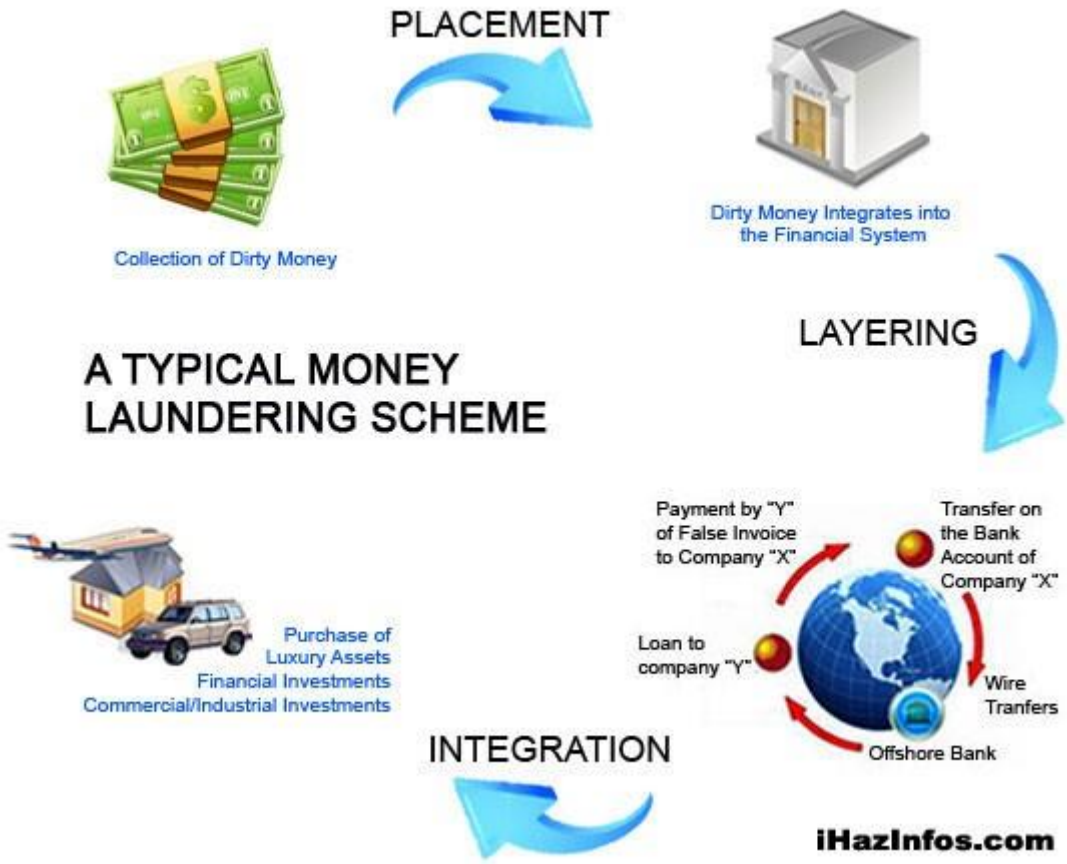

Source: Smooth Criminals; the Hypocrisy of the Drug War, Sergio Casesmeiro, TalkingDrugs,