

Brought to you by:



Artificial Intelligence & Cybersecurity

for
dummies[®]
A Wiley Brand

Understand
cybersecurity

Fathom artificial
intelligence (AI)

Apply AI
to cybersecurity



Ted Coombs

IBM Limited Edition



Artificial Intelligence & Cybersecurity

IBM Limited Edition

by Ted Coombs

**for
dummies**[®]
A Wiley Brand

Artificial Intelligence & Cybersecurity For Dummies®, IBM Limited Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2018 by John Wiley & Sons, Inc.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. IBM and the IBM logo are registered trademarks of International Business Machines Corporation. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN: 978-1-119-50825-0 (pbk); ISBN: 978-1-119-50829-8 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Project Editor: Carrie A. Burchfield

Editorial Manager: Rev Mengle

Acquisitions Editor: Steve Hayes

Business Development

Representative: Sue Blessing

Production Editor:

Magesh Elangovan

Table of Contents

INTRODUCTION	1
About This Book	1
Icons Used in This Book.....	2
Beyond the Book.....	2
CHAPTER 1: Understanding Cybersecurity	3
Looking at the Various Aspects of Cybersecurity	4
Social engineering and phishing	4
Introducing ransomware	6
Malware intrusion.....	7
Non-malware intrusion	8
Detect, Respond, and Mitigate.....	11
Responding to and Recovering From Cyberattacks and Security Events	13
Meeting the Challenges of Cybersecurity.....	13
CHAPTER 2: Fathoming Artificial Intelligence	15
Seeing the Big Picture	15
Teaching Machines to be Smarter.....	16
Learning Algorithms.....	17
Supervised learning.....	17
Unsupervised learning.....	18
Being Smarter	18
Interacting with Humans.....	19
Natural Language Processing.....	19
CHAPTER 3: Discovering Machine Learning and Deep Learning	21
Deep Learning and Deeply Layered Neural Networks	22
Deep Blue plays chess.....	23
Looking at the <i>Jeopardy!</i> champion.....	24
Introducing cognitive computing.....	25
Structured and Unstructured Data	26

CHAPTER 4:	Applying Machine Learning and Deep Learning to Cybersecurity	29
	Predictive Analytics	30
	Taught Not Programmed	30
	Uncovering the needle in the haystack.....	31
	Introducing cognitive computing.....	32
	Identifying root cause.....	33
	A Smarter Adversary	34
CHAPTER 5:	Using the Cognitive Capabilities of Watson to Investigate Security Incidents	35
	Taking Intelligent Action	35
	Understand, Reason, and Learn.....	36
	Applying Watson and Qradar	37
	Winning with Threat Intelligence.....	38
CHAPTER 6:	Ten Trends in Cybersecurity	39
	Responding to Ransomware.....	39
	Combining Application Development and Cybersecurity	40
	Using Deep Learning to Detect DGA-Generated Domains	40
	Detecting Non-Malware Threats	41
	Adaptive Honeypots and Honeytokens.....	41
	Gaining a Better Understanding of How Neural Networks Work.....	42
	Employing Capsule Networks	42
	Deep Reinforcement Learning.....	43
	Protecting the IoT.....	44
	Predicting the Future	44

Introduction

Artificial intelligence (AI) today is a world-changing tool that enhances humankind's abilities in a number of areas. We face the terrifying reality of what malicious hackers have done to our privacy and the grip of terror they keep us in. Never knowing when they'll use their evil prowess to destroy our credit or shut down the local power plant, what we need now are smart machines that help security professionals in the fight against this evil reality.

Computer and network intrusions have shut down airports and hospitals, interrupted commerce, and held people and businesses ransom for their data. The more data that's stolen the more power evil doers have to create exploits that trick you into the foolish behavior of clicking malicious links.

Cybersecurity is one of the greatest challenges of our generation. It seeks to protect our world's data, ideas, and processes; thwart criminal enterprises that prey on our businesses; and exploit people around the world. This area is one of the most understaffed industries in the world with unfilled cybersecurity positions to be about 1.5 million by the year 2020. Help is needed to make the current security professionals more efficient and augment their intelligence. This assistance is coming from AI.

About This Book

Both AI and cybersecurity are broad and poorly understood fields. This book helps give you an overview of the various technologies that make up AI, where they have come from, and what AI has evolved into today. Cybersecurity is another field that has evolved over the last few decades. It's no longer about fending off hackers and quarantining viruses. Dive into the world of cybersecurity and then learn how AI is being applied to the battle. When you're done reading this book, you will be spouting terms like *cognitive computing*, *machine learning*, and *deep learning*, and know how they apply to the cybersecurity space. Moreover, you will understand the important trends in both fields and how AI is the future of cybersecurity.

Icons Used in This Book

You find the following icons sprinkled in the margins throughout the book. They point out useful and helpful information.



TIP

Tips point out important things to know and suggest actions you can take.



REMEMBER

These are the ideas or concepts worth remembering. Don't forget to read them.



TECHNICAL
STUFF

While this book is written for everyone, there are a few places where you might find it interesting to go a bit deeper.



WARNING

Read these to find things you may want to do or not do to avoid something negative.

Beyond the Book

This is a small book that touches on two really large topics, AI and cybersecurity. To learn more about how IBM is using AI to revolutionize cybersecurity, visit www.ibm.com/security.

IN THIS CHAPTER

- » Looking at various aspects of cybersecurity
- » Understanding the cycles to the modern incident response
- » Recovering from cyberattacks and security events
- » Meeting the challenges of cybersecurity

Chapter **1**

Understanding Cybersecurity

Cybersecurity is your only defense in one of the longest wars the world has ever known. Battles are fought daily against nation states, organized crime, thieves, terrorists, and bored, but smart, kids. This war escalates every day as the battlefield grows. Code that can be exploited is everywhere, in watches and phones, smart bulbs and smart switches, thermostats and nuclear reactors.

The cybersecurity defense that has been mounted is staggering. According to CSO from IDG, the expected cybersecurity budget for 2021 will be \$6 trillion. It is not possible to create a perfect defensive barrier against everyone who might want to access computer systems that don't belong to them. The human element in computing has proven to be one of the weakest links in creating a defense. This chapter gives you an overview of the diverse parts of cybersecurity and an introduction to threat response.

Looking at the Various Aspects of Cybersecurity

Cybersecurity is multifaceted in the same way you might provide security for your own home — lights, an alarm, video surveillance, and locks. Protecting an enterprise is far more complex. The job of protecting enterprises has changed over time to meet the demands of a changing threat landscape. With each advancement in cybersecurity, an equal or greater advance is made by those attempting malicious access.

Security professionals have realized that a bad actor with enough sophistication, time, and money wanting to breach their networks will ultimately succeed. In protecting a network and its endpoints, it's what you do next after an intrusion and how quickly you do it that matters.

Social engineering and phishing

The problem with networks, clouds, computers, and connected devices is that people use them. It would be much easier to protect all these things if you didn't have to introduce people into the equation.



TIP

Educating the people that use your network against those trying to trick their way in could go a long way to improving security.

Social engineering is a confidence scam that convinces unsuspecting people to provide information to bad people trying to steal their information. Social engineering is usually the first step in an otherwise complex attack. Think of it as opening the door to someone you believe is a trusted friend.



WARNING

You can easily underestimate the impact of social engineering on data security because it doesn't feel technical enough or software-based. Between 66 and 84 percent of all network intrusion contains a social engineering factor.

The number of large releases of personal information through hacks over the last few years means that almost everyone who has ever used a computer has had some aspect of their personal information stolen. Quite simply, bad people know something about you. Add to this, the fact that “privacy is dead” due to the amount of legal information collected by marketers, social networks, and

all the other ways information is collected, sold, and shared about everyone. A big issue is that social networks have made sharing your personal information an art. We love sharing our doggie pics, where we eat, where we travel, and the names of all our friends and family. How many people use their dog's name in a password? It's just ripe with information. It's not a big task for a well-financed team to correlate all that data with the end-result being that malicious people know enough about you now to be experts at conning you.

Phishing is the further attempt to collect sensitive information about you that may not be readily available online. There is an endless list of phishing examples, but they tend to fall into specific categories usually involving an urgent request:

- » **Renew your subscription or membership:** This preys on your underlying belief that things/routines get disrupted if you don't respond and renew ASAP.
- » **Government record updates:** You may experience a fast heart beat when you see an email from the IRS, Social Security, or Homeland Security asking you to provide updated information.
- » **The embedded link:** Click this link to see your account information, or to learn more, or to avoid a penalty. Clicking that link will most often open malware files that unlock access to secure systems.
- » **Locked and suspended accounts:** These phishing emails usually appear as though they're from a financial institution. One popular phishing email appears to be from PayPal (because many people use PayPal). But with sophisticated data mining, it's also likely that the email will appear to be from the victim's actual bank.
- » **Greed-based:** You've got a refund. Bank error in your favor. You won two free tickets. These and other financial reward emails ask you to provide banking information for direct deposit, your social security number for verification, or many other types of private information in order to release the funds to you — which of course, never happens.
- » **Delivery notifications:** Email that appears to be from a delivery company can prompt you for further information, such as a phone number, home address, or once again, provide an embedded link to malware.

While phishing has been around for a long time, a more targeted attack based on research specific to the victim called *spear phishing* has become popular in recent times. Email, social media messages, or other information can be sent to you, and you believe the messages are from a trusted source such as an employer, frequently visited website (also known as a watering hole attack), a friend, person you know, or from companies where you regularly do business. The information used to create these kinds of attacks come from stolen information. If you've done business with a retailer that has had its data stolen, the bad guys know where you shop. So, you might receive what appears to be a friendly email warning you that your information has been stolen and to "click this link to verify your account." More data is stolen every day, and that information is being used to create spear phishing attacks.

The larger your enterprise, the higher the chance that someone in the organization will fall victim to social engineering. The answer to protect against social engineering is education and awareness. Teach employees to look at the URL of origin, search websites about common scams, or even call the organization on the phone to verify the request before providing information.

Introducing ransomware

To understand ransomware, you have to have a basic knowledge of something known as *public key cryptography*. Public key cryptography has given rise to some of the most amazing advances on the Internet, from being able to safely communicate with a bank and securely sending credit card information to communicating privately between peers. With this type of encryption, public keys are shared and private keys are kept close to the vest. Keys are simply long strings of numbers and letters. Hand out your public keys to whomever you wish so people can encrypt messages meant only for you. Messages encrypted with your public key can only be decrypted with your private key. When text is encrypted with a private key, only copies of the corresponding public key can decrypt it. This is what is commonly known as a *digital signature* because decrypting something signed with your private key proves you're the one that signed it (encrypted it).



REMEMBER

The most important thing to remember about public key cryptography is that it's only as good as your ability to keep private keys safe. So, if your computer is hacked and you have a file in your documents folder labeled MYPRIVATEKEY, chances are good that any privacy you once hoped you'd have from using public key

cryptography is now gone. The hacker can now also digitally sign things making people believe they came from you.

For all its good, public key cryptography was used to create something known as the *cryptovirus*. This virus infects a computer and uses public key cryptography to encrypt the files on the computer that can only be decrypted with a private key. The hacker then holds your files ransom and asks to be paid some amount to decrypt your files. Thus the name *ransomware*.

Here again, public key cryptography is in play. This type of encryption is the foundation of what we now know as cryptocurrencies like Bitcoin or Ethereum among the many others. One of the basic elements of any good cryptocurrency is that its transactions are anonymous. This creates a perfect way for malicious people to demand payment without fear of getting caught. They supply, for example, a Bitcoin wallet address that can't be traced to them and demand that you send funds to the wallet.

One form of ransomware is called *cryptolocker* and infects primarily PCs. It is spread by hiding malware in things such as Microsoft Word macros and then distributing a Word document that you try to get people to open using social engineering. After a Windows PC is infected, it spreads the virus to other operating system types such as Linux and macOS. One of the worst of these is known as the *Locky virus*. Starting in 2016, the virus infected hundreds of thousands of computers. At least 23 million infected email messages were sent as of August 2017.

Ransomware isn't just for individuals. In 2017, the WannaCry ransomware virus infected the British National Health Service. In this case, the virus exploited unpatched versions of Windows and quickly took down many systems, causing appointments to be canceled and operations to be rescheduled.



REMEMBER

Back up your data. While most modern enterprises have strict backup policies, the average individual does not. Modern enterprise endpoints are all at risk. Backups are the answer to not being worried about having your data encrypted and held for ransom. The challenge then becomes how quickly you can clean the computers and restore your data.

Malware intrusion

Roughly six kinds of malware can infect computer systems. It's believed that about half of these are *computer viruses*. These are

small programs that once they infect a computer replicate themselves and attempt to send the replicated version to another computer, much in the same way a biological virus infects its host. Most of the virus programs impact the computers on which they find themselves by destroying or altering data.

The computer *worm* is a type of malware, which like a virus, does its best to replicate and spread itself. Unlike viruses that need a host program in which to be spread, worms are standalone programs that can spread completely on their own. They do the same kind of damage as a virus.

The *Trojan* is malware named after a mythical wooden horse used by the Trojans more than 2,700 years ago as a peace offering to Sparta. Once the Spartans brought the Trojan horse within the walls of Sparta, soldiers hidden inside were able to attack and defeat Sparta.

Trojan computer malware is often disguised as legitimate software programs. Any number of social engineering methods are used to get people to download and install the Trojan. Once installed, it usually works to give access to a malicious third-party. Unlike viruses, they don't replicate themselves. They rely on the con game to get people to infect their own computers.



REMEMBER

The weak link in even the most protected network is the people that use it.

Similar to the Trojan that allows malicious access to computer systems, the *Trojan downloader* has a single purpose: downloading malicious programs from a remote server. The downloaded programs are most often used to create *botnets*. This is a network of infected computers controlled by a central computer. These are often used to conduct *distributed denial of service (DDoS)* against websites and Internet infrastructure. When thousands of infected computers begin hammering an IP address or block of IP addresses, they quickly make it impossible for anyone else to communicate with the web service.

Non-malware intrusion

Non-malware intrusion involves security breaches not caused by a program that has breached your defenses. I know it is probably improper to define something based on what it isn't. But with non-malware intrusion that's the easiest way to encompass such a broad arena of attack vectors. It generally doesn't involve

a virus, Trojan, or other malware file. Generally, a non-malware attack is an attempt to breach a computer or network using software you trust, like Microsoft Office or the Windows Management Instrumentation (WMI). Even though the first two examples were Windows-based the software commonly used in these types of attacks can run on many other operating systems, such as Linux (and all its variants).

SSH, the Secure Socket Shell, is a hacker's dream. It's the administrator's "back door." It provides remote command-level access to whatever operating system is running SSH. Today, that would be all of them. Windows was the last holdout to add SSH access.

You've all seen the movies of the super hackers who sit down at a keyboard and just before the bomb timer ticks to zero they've gained access to the remote network and shut down the bomb. The most common way into an SSH account is to guess the password. That doesn't sound very efficient, but with modern brute force password-guessing tools the job has become much easier and more reliable.



TECHNICAL
STUFF

The SSH brute force attack has two simple elements:

- » **Find an SSH server to attack.** Use NMAP or a port scanner to find an open SSH server.
- » **Use a brute force tool.** There are several brute force tools like Hydra, NCrack, and Medusa that use massive word lists to guess passwords.

A few things you can do to attempt to fight off SSH attacks include

- » **Running SSH on a non-standard port, the standard port being 22:** That will keep out the complete novice hacker.
- » **Blocking the SSH login for the root user:** That makes gaining access significantly more difficult because now the program also needs a username to go along with that brute force password finder. That is why you must train your users not to fall prey to phishing attacks. A user might think, "All they want is my username but not my password; this must be safe." Nope!
- » **Limiting the number of tries a user can make to guess the password until the account is locked:** This keeps many of the brute force programs from making endless guesses.

Structured Query Language (SQL) is the query language for many modern databases. Malicious instructions can be embedded into an SQL database, causing it to return the contents of the entire database (commonly referred to as an *SQL injection*). This can be done by entering the code into a user input field in a web form. This code then gets embedded as a user field in the SQL sent to the database.

This is an old exploit, and today's enterprises have updated their database query codes to eliminate this kind of attack. One method is using something known as *parameterized SQL statements*. This pulls the user input information out of the actual SQL statement itself and places it into parameters that are passed along. This kind of programming has largely defeated the SQL injection attack. But, there are many unpatched systems out there. It's a bit of work to go through and recode all those SQL calls in your program. Not doing so means your Internet-facing database is just not secure. So, problem solved?

Newer database types have arisen that are allowing for data to be distributed across the Internet, making data access to mobile applications much more efficient. These are called *NoSQL databases*, one popular example being MongoDB. As you might expect, a query language is not used to retrieve data from a NoSQL database. There are different types of NoSQL databases that use various methods of storing and retrieving data, but they have not proven immune to injection attacks. Entire books can be written about injection attacks. Suffice it to say that malicious code can be fed to NoSQL databases like MongoDB or the superfast in-memory database NodeJS. There are tools for checking the data sent to these new databases, but they aren't perfect.

The least technical, but possibly the most difficult to defend against, is straight up credential theft. This is most often done using sophisticated phishing and social engineering. Once the user's username and password has been compromised, it's all over. For the well-funded or state sponsored cyberattacker, it's also possible to send human beings to look for handwritten username password sticky notes. Imagine the janitor each night photographing the sticky notes found on the front of monitors or on desktops everywhere. Better password management and stricter password requirements have had both a positive and negative impact.

Forcing users to have longer and more complicated passwords and changing them regularly has also forced them to write them down. The alternative to writing down the password is using one of the new secure password managers. Yet, some of the password management services have been hacked, and you can never be really certain that you aren't using a malicious clone program. Two-factor authentication has gotten some of this under control. This works by requesting a code sent to an app on your phone or to your email before allowing access to untrusted computers. But, the password is largely dead, and there is now a huge move to implement better biometrics like facial recognition, voice printing, fingerprinting, skin sensitivity, and heart sounds, or combinations of these.

Detect, Respond, and Mitigate

While many network and computer intrusions go undetected, the job of a security professional is to figure out when an incursion has happened so something can be done. It's the "uh-oh" moment.

There are different ways to monitor a network for intrusion. One of those is *anomaly-based detection*. This is like walking into a room and knowing something is out of place. At first it might not be obvious what it is, but every fiber of your being knows that something is wrong. Looking at the landscape of network activity is the same. There are system monitoring programs that alert you when something is out of the ordinary. This might be a change in the level of data access that may be a sign of a Denial of Service (DoS) attack. It may be a change in the way a person is using her computer that alerts the security team to a possible malware intrusion.



WARNING

One of the significant problems with anomaly-based detection is the high rate of false positives. Following each of these anomalies is time consuming, and yeah, ultimately expensive.

One of the other common ways to be alerted to network intrusion is the use of deception to set a trap. Deception in nature is pretty interesting. It is one of the fundamental ways organisms have protected themselves over the millennia. One type of detection says, "Nope, not me. You've got the wrong guy. Eat someone

else.” This is done by mimicry, making yourself look like something not as appetizing as you otherwise might be. The other is completely the opposite. It says, “Come here. Tastiest meal ever.” Yes, that’s right, the mousetrap.

In cybersecurity, digital honey is used as bait. The traps, known as *honeypots* and *honeytokens*, are designed to lure intruders away from the good stuff. Honeypots are computer systems on a network that appear to be full of data, just waiting to be stolen, but in fact contain only the bait. Normal network users never use the honeypot. So, when someone tries to gain access, that person triggers an alert.

Various kinds of honeypots increase the chance that someone will fall for the bait:

- » A *low-interaction honeypot* is normally an emulated network service running on a special computer, for example, email, ftp, and web servers. These fake services tend to be invulnerable, thus protecting the honeypot computer. Essentially, they are built from an open port and a listening socket and nothing more. But they do trigger an alarm. They are lightweight, and many of them can be run simultaneously on a single computer. Sadly, it’s like putting a fake owl in the garden to scare away the birds. Soon, the birds are sitting on the head of the plastic owl. They don’t fool people for very long.
- » A *high-interaction honeypot* is the real deal. It’s a computer system running on the network with real services providing fake data. These computers have network monitors that begin tracking intrusions and exploits from anyone trying to interact with them. As you can imagine, the real limitation here is the amount of effort it takes to maintain this kind of honeypot. Again, there is the problem of weeding out the false positives.

In today’s world, data is distributed across networks, mobile devices, the cloud, and Internet of Things (IoT) devices. Once an intrusion occurs, it becomes important to track where the attack originated. Fake data, email addresses, and fake accounts known as *honeytokens* do just that. This fake data is seeded across the network, and a record is kept of where it was placed. When data

containing a honeypot is stolen, you know where it was stolen from. The same thing happens when honeypot email addresses or accounts are used. Cartographers do this same thing when they add a fake street to a map just to prove that someone copied their map. This is a good segue into what to do next, respond.

Responding to and Recovering From Cyberattacks and Security Events

Once an intrusion attempt or an actual intrusion has been detected, it's all about investigation and response. How quickly can a security team identify the intrusion, malware, or non-malware; determine if it's a false positive; figure out the method of intrusion; learn as much as it can about the intrusion; close the door that was used; and then get rid of the problem?



REMEMBER

The best cybersecurity defense includes well-trained users. Their response when they suspect something might be wrong is critical. Train them to be paranoid.

According to a 2017 Ponemon Institute study, the time it takes to respond to a security event was lowered, “from an average of approximately 201 in 2016 to 191 days and the average days to contain the data breach from 70 to 66 days.” While this is headed in the right direction, the complexity of protecting and responding to information environments spread across everything from smart devices to cloud hosting only becomes more difficult. Add to this complexity response across a multinational organization and you have a costly nightmare on your hands.

Meeting the Challenges of Cybersecurity

Smarter software is the trend in cybersecurity. Security Information and Event Management (SIEM) software provides analysis of security events and the storage and correlation of a wide variety of information, including such things as log data, threat vector and user behavior, and analysis of structured threat intelligence.

In addition to threat response, there is a move to make software application development more secure as well. Integrating cybersecurity into the phases of software development helps manage the risk involved in releasing software that will immediately be attacked by hackers looking for mistakes and vulnerabilities. Code is everywhere. It's tempting to downplay the software running on a smart refrigerator, but if it sits on a local area network, it's a point of vulnerability and its code is just as important as any productivity application. Nothing is overlooked by anyone wanting to pierce the network. Nothing should be overlooked by those wishing to protect it. The result of a network intrusion can have a huge impact on organizations large and small.

IN THIS CHAPTER

- » Getting to know pattern recognition
- » Making machines smarter
- » Moving beyond human intelligence
- » Accessing the world's information
- » Talking to smart machines

Chapter 2

Fathoming Artificial Intelligence

Artificial intelligence (AI) is the attempt by humans to make machines smart. Intelligence is such an important part of what makes humans unique that until recently the embodiment of intelligence in a machine was almost always accompanied by a humanoid robot. Today, along with the development of smart devices that assist us, one of the important goals of AI is to provide answers we would never arrive at by finding those answers within our own massive archives of information.

AI may prove to be the largest advancement in human technology since the start of the industrial revolution. Once conceived only in science fiction, AI is finally here and impacting daily life.

Seeing the Big Picture

One of the things smart machines are really good at is analyzing data, such as text and images, by using a process known as *pattern recognition*, considered a branch of machine learning. Pattern recognition uses both supervised (training data) and unsupervised (no training data) to find patterns in data, either visual or textual.

Most visual pattern recognition is done using supervised learning algorithms. A significant number of training images are provided for the computer to learn and be able to recognize a pattern. Pattern recognition in text data is sometimes called *data mining*. One example of its use today is the Gmail auto response system that makes suggestions of email responses based on the content in your email. Spooky, but it works.

We live in a fuzzy world of imprecision. If you grabbed your reading glasses to read this, you really understand this concept. Most pattern recognition uses a type of *fuzzy logic* where answers are close enough but not exact matches. Even the way AI performs math calculations allows it to arrive at answers that are “close enough.” Not that a computer really says this, but if you asked an AI system to add $1 + 1$, the answer might be “Well, based on other addition examples I’ve been given, it’s close enough to the value of 2 to be 2. So, I’m saying 2.” When mining data, computers are often confronted with subjective adjectives like young and old. The machine learning algorithm must be able to handle this kind of imprecision.

Teaching Machines to be Smarter

Some great futurist minds have discussed a point in time they call *the singularity*. This is the theoretical future date when a super intellect is born (made) that is smarter than humans in nearly every field. This will be the age of machines, replacing the age of man according to this theory. This eventuality does not come without its naysayers. There are some that are worried about what a super intellect machine might see as necessary for its own, and Earth’s, survival. The goal then becomes creating a *friendly AI*. This is a form of *Artificial General Intelligence* that has a set of ethics that it follows for the betterment of mankind.

Artificial General Intelligence (AGI) is the goal of having machines at least as smart as humans, if not smarter. Significant advancements in computer hardware are going to be necessary to achieve this. Human brains, described in computer terms, can perform 38 thousand trillion operations a second. Not even our fastest supercomputers can come anywhere close to that. So, it’s going to be a while before we have a super intellect in a single machine.

There are many steps along the path to having a single machine that equals or surpasses human intelligence. One of those steps, known as *narrow AI*, is having AI computers specialize in a particular field of study. They don't really have a world view beyond solving a particular task. For instance, chess playing computers eventually beat even the best human grand masters. This is a type of AI sometimes called *reactive AI*. The reactive part comes from the fact that the computer is responding to challenges in the present. It has a goal, looks at what might help it reach that goal, and makes the next move.

In many respects, distributed machines, connected by a network, are already smarter than humans. Humans have limitations. It's not always clear what those limitations are, and they vary between individuals, but the limitations largely have to do with input. How much information can a human have access to at any one time? A human's ability to make sense of large data sets, compared to that of AI, is not very good.

Learning Algorithms

There is no single way to design machines that learn. The underlying code contains *learning algorithms* that are programs that extrapolate insights (intelligence) based on data provided to the computer. There are two basic categories of learning algorithms, supervised and unsupervised.

Supervised learning

Supervised learning is exactly what it sounds like. Someone supervises the input of information upon which the learning algorithm will arrive at a conclusion. Think of this like giving the computer a tutor.

One of the most basic supervised learning algorithms is designed around a decision tree. This is the foundation of the expert system, a series of yes and no questions sufficient for the computer to arrive at some probable answer. With an expert system, a conclusion is derived based on the programmed inputs of field experts. For example, diagnosing starter problems in a car will require the user to answer questions about the symptoms experienced when trying to start the car. Do you hear a click when you turn the key?

Yes or No. Based on that answer, new questions along the tree are asked until the computer suggests, “Your battery is likely dead.”

More advanced types of learning come from the use of training data. Unlike an expert system where specific answers are provided by experts, allowing a computer to learn by training provides unique capabilities. Feed the computer hundreds of thousands of cat photos and eventually it will be able to recognize a cat in a photo. This is a step up in learning because it’s based on probability. Based on every other cat picture the computer has seen, it forms a probable idea of cats in a photograph. Two common types of learning algorithms of this type are logistic regression and a back propagation neural network.

Unsupervised learning

Unsupervised learning allows for the training of AI, using data that’s unlabeled and unclassified with the use of special algorithms that allow the AI to learn on its own rather than being spoon fed the data by a human. Two common unsupervised algorithms include the *apriori* and the *k-means*.

Being Smarter

Many decades have been spent trying to make machines smarter. Smarter comes not only from new and better software but also advances in machine technology. Computers run significantly faster than they did 20 or 30 years ago. There have been advances in sensors, cameras, and microphones that have given computers an improved ability to take in its surroundings.

Incredible databases have been compiled with machine intelligence. For example, the Open Mind Common Sense project at MIT has collected more than a million facts since its start in 1999. This particular database contains statements of fact that we might consider common sense. “Fire is hot.” These short sentences are later parsed into structured data so they can be easily understood and used by a computer in the same way it might search a library catalog for a book author based on the title of the book.

Inference can then be drawn from all the pre-programmed facts. For example, if fire is hot, and logs can catch fire, a burning log is also hot. The ability to draw inference from a large number of facts leads to a very smart computer. Careers have been spent collecting data and structuring it. The only problem is that most of the world's information is unstructured (unlabeled, unexplained, uncategorized, and embedded within text).

Interacting with Humans

There is a branch of AI working on simulated human emotions, primarily so robots can be used as human companions. Computer systems have something known as an I/O system (input/output). Humans really aren't different in that respect. The challenge is to get the two I/O systems communicating with one another. Earliest attempts at interacting with machines came from using the only electronic keyboard of its day, the 19th century teletype. These keyboards were used to create punch cards used by 1930s IBM adding machines. Eighty years later, keyboards remain the predominant mode of communicating with machines. Output from the machines was similar. Responses were typed on paper using a teletype and eventually onto a cathode ray tube (old style computer monitor).

What everyone really wants is a computer to have a conversation with. It's not a new idea. The idea of a talking machine goes back nearly 200 years when Christian Kratzenstein built the first speaking machine in 1773.

Natural Language Processing

In 1954, a machine was used to translate 60 Russian sentences into English using an IBM 701 mainframe computer. This success is largely seen as the beginning of modern natural language processing (NLP). While this experiment used punch cards, the goal of NLP is to have computers understand human language as it is spoken.

NLP is the heart of the modern AI input/output system. The ability to understand natural language provides two important capabilities. The first is the ability to comprehend unstructured data, such as the 2.5 million peer reviewed papers published every year. The second is perhaps the one that the general public is most familiar with, the ability to talk to a computer and have it understand our natural language requests.

Even though NLP is in its early stages, home smartspeakers are quickly becoming pervasive. They can tell you what time it is, the weather anywhere in the world, convert measurements, do web queries, help you shop, play games with you, or play your favorite music.

While all knowledge is not embodied in a single machine, or even shared among clustered machines, the fact remains that computers can now assist us by doing the things we aren't that good at. They are tireless, providing emotionally unbiased responses by correlating the sum of human knowledge on a particular topic. What might take a human weeks or months of research can be presented in a matter of moments.

- » Understanding self-learning computers
- » Beating Garry Kasparov
- » Winning at *Jeopardy!*
- » Seeing how cognitive computing works
- » Analyzing the world's data

Chapter 3

Discovering Machine Learning and Deep Learning

Computers have come a long way from simply being data storage repositories with a few favorite applications that allow us to view and crunch that data. The road to machines that think has been long and winding. In the early days, databases with query languages allowed humans to construct evermore clever queries in order to arrive at insights from the data contained within.

Machines became more powerful when programmers created algorithms (fancy word for program snippets) that would process data based on rules derived from insights. Chess computers, armed with the rules of the game, were also given algorithms that told them which strategy might help them win the game given a particular layout of the game pieces.

The big leap forward from smart gaming computers that relied primarily on high-speed, brute-force examination of large sums of data, were computers that understood language. This allowed computers to make sense of human-written text and have a conversation about it. This not only made computers seem smarter, but also it made them more intelligent.

While natural language processing (NLP) powers front ends to things like switches, lights, and thermostats, it has also allowed for a new kind of learning, where machines don't just collect data; they get smarter, and those smarts are used to augment human intelligence.

Deep Learning and Deeply Layered Neural Networks

Deep learning is a branch of artificial intelligence (AI) and a subset of machine learning that allows computers to learn on their own, unsupervised, using neural networks. The cool thing is that neural nets work in a fashion similar to the human brain. No, it doesn't need coffee in the morning!

What deep learning does is read unstructured data and form patterns and clusters of patterns. This is similar to how the neurons in the human brain work to form thoughts. At least that's the theory put forward by researchers McCulloch and Pitts in 1943. Their model of the neuron and how it works is still used in AI today. Figure 3-1 gives you a simple example of a neural network. This figure shows how layer 1 accepts input, how outputs are summed and sent to hidden layers where transformation functions do their stuff, and then how those outputs are summed again into an output layer. Think of it like a drip coffee maker. You mix the inputs, hot water and coffee grounds, filter the output of this mixing, and then mix it all together into the stream that ends up in the carafe.

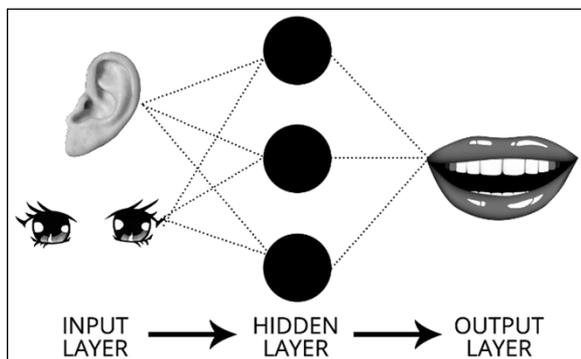


FIGURE 3-1: A neural network.

Deep Blue plays chess

In 1985, a group of researchers at Carnegie Mellon University created a chess playing computer called ChipTest that was capable of searching through 50,000 moves per second. The team was successful enough to be hired by IBM Research to develop the next generation chess-playing computer called Deep Thought, after the fictional computer in Douglas Adam's *Hitchhiker's Guide to the Galaxy*.

Deep Thought was the first in the line of chess computers to beat a grandmaster when it won against Bent Larsen in 1988. The following year it lost to the famed chess grandmaster, Garry Kasparov. The computer was renamed in 1989 to Deep Blue. (IBM was often called "Big Blue" based on the color of its logo.) After several upgrades, Deep Blue became the first computer to win against a reigning world champion when it beat Garry Kasparov in 1996.

Was Deep Blue smarter than Garry Kasparov? Deep Blue won using sheer brute force. Current versions of Deep Blue can evaluate more than 100,000 chess moves a second, where Garry can evaluate about 3 moves a second.

So why build a chess-playing computer? A machine capable of evaluating the results of a chess move, after analyzing hundreds of thousands of possible moves, was the foundation of the ability to arrive at predictive values (chess moves in this case) after analyzing both the "rules of the game" and the results of numerous games played by grandmasters. This same idea can be applied to almost any field. For example, in chemistry, there are physical rules about how molecules interact. Given a set of rules, the computer could predict the result of various chemical compositions? Is this AI? Not really. Instead, it's the result of supercomputing, or really fast computing originally measured in something known as *gigaflops*. One gigaflop equals one billion math operations a second. In case you're interested, the fastest computer in the world as of 2016 is a Chinese computer that can perform 93 quadrillion operations a second (93 petaflops).

While Deep Blue didn't use what we might currently consider AI, it was the birth of a belief that a computer could beat a human at a game that requires strategy and skill. There was no deep learning in Deep Blue. Every move had to be calculated in the moment.

That's why for every increase in processing power came a corresponding increase in the ability to win chess games. The bottom line was that Deep Blue proved we could make computers do something better than a human.

Looking at the *Jeopardy!* champion

The television game show competitor, Ken Jennings, made history in 2004 by winning 74 consecutive games of *Jeopardy!* In the middle of this Guinness record streak, IBM executives had the idea that it would be interesting to have their computer, Watson, compete on *Jeopardy!* (Watson was named for IBM's first CEO, Thomas Watson.)

Training Watson was slow at first. It meant feeding answers and the corresponding questions to the computer manually. Large teams of student volunteers helped. These answer-question pairs came primarily from previous *Jeopardy!* shows that had never aired. By the end of 2006, Watson was only able to answer questions correctly 15 percent of the time. By 2010, Watson was winning in competitions against *Jeopardy!* contestants regularly.

The broadcast competitions of Watson versus two other *Jeopardy!* champions, Ken Jennings and Brad Rutter, were finally aired February 14th and 15th, 2011. The intelligence displayed by Watson during these games went much further than simply being able to answer questions correctly. It also displayed a complex strategy by betting based on Watson's confidence level within the game category. Watson ended up winning the \$1 million prize, which was then donated to charity.

The advance displayed by Watson over its chess playing predecessors was its ability to communicate using NLP. Watson needed to be able to parse the answer and correctly form the corresponding question. It also needed to be able to analyze the questions and answers from previous games with the answers certainly worded differently.

Winning at *Jeopardy!* was just the beginning for Watson. Since that time, Watson has applied its learning and interactive capabilities in many fields that include everything from recipes to cancer research and even more recently tax return preparation and of course, cybersecurity.

Introducing cognitive computing

AI is the ability, displayed by machines, to perform tasks that would require intelligence if those same tasks were performed by humans. How those tasks are actually accomplished has changed over time and so has the idea of the very nature of AI. The idea of AI has grown to encompass more fields of study and new ways to accomplish intelligence-requiring tasks.



REMEMBER

AI has advanced from expert systems arriving at answers by accepting yes-no questions, through massive and high-speed data analysis required to beat grand champions at chess, to being able to communicate using natural human language in order to both interact with humans and analyze text data.

Cognitive computing is the next step in AI. Previously, AI computers could create a log of their responses and even how they arrived at those responses. But they didn't really learn anything having arrived at a response. It wasn't one bit smarter having done the research and answered the question. If you ask the same question 15 times, it would give you the same answer 15 times.

It's easy to get enamored with the capabilities of AI and forget that the ultimate goal of AI is to build tools that augment human abilities. Cognitive computing accomplishes that in a different way than simply spewing out answers. Instead, it provides insights based on its research. It also updates that insight when new information becomes available. It may ultimately arrive at insights that a human would not.

Cognitive computing simulates the thought process of a human brain. It learns, collects information (data mining), understands patterns, and is able to communicate using human language. More importantly, cognitive computing fills a slightly different niche than its counterparts. It takes a back seat, as it were. Imagine the difference between a DMV driving tester saying, "Turn right at the next corner," and someone sitting in the back seat talking in your ear saying, "There's traffic if we go right even if it's shorter. Turning left will take a bit longer but involves fewer lights, and there is a coffee shop along the way. That might be a more enjoyable route." In this example, the human is still the driver and

makes decisions based on the insights provided by the computer. It can then be said that cognitive computing uses a “humans-first” approach.

With machines getting smarter and faster at a dizzying pace, this back-seat companion approach is a much less fearful way forward. Some of the smartest people around the world were suggesting that AI might end up destroying humanity. This new approach steers (literally in some cases) away from that scary scenario.

Structured and Unstructured Data

Structured data, tagged in some way to make it easily understood by a machine, will continue to provide massive amounts of data that can be mined for insight. One example of this is the data returned from self-driving cars. Today they are fairly rare, and tomorrow they will be everywhere. Mining this data and the vast sums of data from IoT (Internet of Things) devices like smart watches allows computers to find patterns and, in those patterns, insights into how those devices can behave differently in the future. Each vehicle on the road can provide data to better the abilities of every other car on the road.

Because we know where this data came from, what part of the car, for instance, it’s labeled data. We know, for example, it was a particular sensor, such as battery heat, sending the data. Correlate all that data from every other kind of sensor including GPS location, battery level, speed of the car, and even who is driving, and you begin to see the kinds of insight that can be derived.

Unstructured data tends to be the result of human intelligence applied in a human readable written format. Think of it as pre-digested intelligence. A human brain first processed the data and arrived at insights and then recorded those insights and, in some cases, the data that led to the insight.

A good researcher rarely speaks in absolutes. Instead, researchers will generally present their confidence levels of the insights they reached as a result of the work they did. This is what cognitive computing does. It arrives at one or more, usually more, insights by reading text, examining the work of others, making sense of it, and then proposing conclusions, each accompanied with a confidence level.



Cognitive search and discovery uses AI in the form of NLP, pattern recognition, and machine learning to understand and organize information it collects from a body of knowledge. The body (corpus) of knowledge isn't just huge; it's growing both structured and unstructured knowledge bases every second of every day. The conclusions or insights, and their corresponding confidence levels, are updated on-the-fly, as each new bit of information is pulled in and correlated.

While humans generally have their own corpus of knowledge in which keyword searches are done to find relevant content, AI systems are plugged into the data using APIs (Application Programming Interfaces).

A great example of a knowledge base that allows computers to connect via an API is WikiData. It's a structured collection of data based on the plain text information in Wikipedia and has an API interface that allows programs to query it and then returns structured content in either XML or JSON format.

An AI system capable of this kind of research works in the same way humans take for granted. It must derive semantic information from the text it reads. The system will extract relationships between entities, derive keywords from the content, and determine sentiment. In brief, sentiment “colors” the content by figuring out whether people are happy or upset about what they are writing. Writing about the closing of the ozone layer because we've reduced hydrocarbons would have a positive sentiment, while an article about rising sea levels would likely have a negative sentiment (unless you live inland and always wished for a seaside home).

NLP breaks sentences into components such as subject, action verb, and direct object. Remember having to do that in grade school? When we read our brains continue to do that processing at super high speed even if we believe we're doing it without thinking. We'd be wrong.

The computer then detects entities such as people, places, things, and geographic features, along with the content's keywords and then ranks them based on importance. There is a complex method for determining importance, but web search spiders do the same

thing by finding words that are placed in the title of the page, headings, bold type face, and used more than once to determine a keyword's importance.

Then the magic happens. A cognitive computer will take what it learns from examining the text and correlate it to what it already knows and form new concepts that are not necessarily discussed in the text it has read. These insights are then weighted and provided, using natural language, to a human user.

IN THIS CHAPTER

- » Determining risk using predictive analysis
- » Understanding, reasoning, and learning
- » Seeing that the bad guys have AI too

Chapter **4**

Applying Machine Learning and Deep Learning to Cybersecurity

Artificial intelligence (AI) applied to cybersecurity provides security professionals with an augmented ability to protect endpoints, data, and networks. By using sophisticated abilities to predict problems based on prior solutions and an ability to use natural language processing (NLP) to analyze unstructured data, unique solutions and detailed insight are provided to the Security Operations Center (SOC) to quickly and cost-effectively stop intrusions or even prevent them before they happen. It's also the only way to protect a network against malicious attackers also using AI.

Predictive Analytics

Cybersecurity professionals use analytics to detect anomalies in network patterns, network traffic, and normal user activities. Exploits are identified by their *signatures* (known patterns of attack). These are the identifying methods that the malware or attacker has used to gain entry into the network. Network analysis software alerts the security team when a signature attack is recognized. That's all well and good for real-time monitoring but it most always means that the deed was done. Cybersecurity has moved on from a complete reaction to activity to one where networks are managed based on risk. Each entity involved in the network's activity is scored based on the risk. You can think of this like having a credit score, which is also a form of *predictive analysis*.

Predictive analytics gives you a look into the future, albeit fuzzy. One approach, which you might call an “on the doorstep” scenario is being able to identify an intrusion without having a prior signature. Machine learning in AI actually learns how to recognize patterns far better than a human. By analyzing all kinds of previous attacks machines have begun to have a “gut feeling” or predictive ability about what might be an attack, even if it doesn't match a previously known signature.

With the network in constant flux, it becomes a superhuman job to determine exactly what a network's normal behavior looks like. There are also malware programs that sit on the network appearing innocuous because the damage (normally data theft) is long term. These are called Advanced Persistent Threats (APTs). They're cleverly designed to be overlooked by network security programs and to remain in place for as long as possible.

Taught Not Programmed

Artificially intelligent machines today are not the sum of their programming as they once were. They analyze great sums of data, the more the better, and find patterns that might have been otherwise unrecognizable. Machine learning may examine millions of math problems and their results and determine, based on a pattern, what the result might be. Applied to cybersecurity the

goal is to examine this network data and apply everything it has previously learned to augment a human-led security team.

The acronym URL, not to be confused with the Uniform Resource Locator of the World Wide Web, is an acronym that stands for the following:

- » **Understand:** Examine the mass of prior research using NLP. This information can be found within videos, books, magazines, journal articles, and yes, even PowerPoint.
- » **Reason:** Provide insights based on analysis that include what type of attack may occur, or may have occurred, and the types of threat entities involved in the attack and their relationships.
- » **Learn:** Up to the millisecond research findings continually add to the corpus of knowledge. New insights are continually created based on new information.

Nothing stands still. It would be nice if there were a “stop” button. But the reality is that security research is made new every moment, and the bad guys trying to break into the network get smarter and more sophisticated. Humans are generally not good at change. We get tied up in what we know. We get tired. We overlook things. We get pet ideas and form pride in the knowledge we’ve accumulated. When something comes along that says, “All those years learning about computer viruses is now completely useless,” it’s like a punch in the gut. But, not for computers. They tirelessly learn, adapt, and form new ideas without ego about what they already know.

Uncovering the needle in the haystack

If there was ever a bad phrase in cybersecurity, it might be “false positive.” It’s the great time waster and money sink in the world of network security. If you want a visual, imagine the SETI researchers, hunting for alien signals. They’re looking for any little blip on the screen that rises just above white noise. On large enterprise networks, there’s a lot of background noise. So, at what point does the abnormal behavior rise to the level of something for which a security researcher needs to apply attention? This attention begins the immediate application of time and money to what might be nothing.

Both humans and machines can cause false positives. How often have you rebooted a machine because some program has gone crazy and started eating up all the memory and CPU cycles in your machine? Or perhaps a user suddenly decides he needs a local copy of that multi-gigabyte database and starts a download, triggering a traffic alarm.

Software can also cause false positives. Programs that have not been tested for security before release can cause a nightmare of protocol violations that would appear to any good security program like an attack but is really just untested software running on your network.

One of the goals of using AI in cybersecurity is to weed out those false positives. Computers, on their own, even really smart ones, fall prey to the false positive. But, using the insights of a computer in partnership with a human in a type of hybrid approach lowers the risk that non-malicious network use will be seen as a false positive. Because the AI is always learning, it can also learn from its human partner. The security researcher examines the alerts provided by the smart security system and determines which are and aren't false positives. Those results are fed back to the AI, making it smarter. In time, the AI will report fewer false positives.

Introducing cognitive computing

Cognitive computing is AI that emulates the actions of a human brain. Applied to cybersecurity it learns and gains the ability to identify threats by investigating security incidents compiled by security events from a variety of inputs. Cognitive computing machines use input from structured and unstructured data, as well as human-machine partnerships that provide human insight. It then augments a cybersecurity professional by providing insights it may never have come up with or suggesting unique solutions, and all much faster than a human is capable.



WARNING

One challenge of any network security team is fatigue. Responding to hundreds of thousands of security events, most of which require no action, is enough to wear down the best and brightest. Think of this like the chime on the door of a store. How long is it until the clerk no longer even hears that chime. One of the reasons it becomes background noise is that the door chime doesn't really tell them who is coming through the door. Is it a lost child? Is it a shopper? Is it a robber? Without context, the chime quickly

fades from importance. Put a little facial recognition on that chime and suddenly you're paying more attention, as information about each person appears on the screen. Correlate shopping data from a frequent shopper card, and then you have a powerful tool. "Hello Mr. Smith, we have a new crop of bananas now available on the endcap of aisle 6."

Cognitive computing systems provide the kind of contextual information needed by a security professional to make faster and smarter decisions with less fatigue. With data feeding in from network security analytics software, and contextual data updating the cybersecurity corpus of knowledge every second, combined with previously determined insights (self-learning), cognitive computing correlates all that data faster and more accurately than any human.



REMEMBER

It's been more than half a decade since an AI computer trounced the smartest human *Jeopardy!* players. Coming up with solutions faster than a human has only improved over time.

Identifying root cause

Responding to continual exploits without finding the root cause can waste a great deal of time. When analyzing an exploit, determine the causes, support it with evidence, find solutions, take notes, and recommend actions. The actions you ultimately take are based on two of these analysis points: discovering the root cause and then finding a solution.

Root cause is sometimes an enterprise systemic thing. Therefore, it makes sense to find the overarching cause rather than pointing a finger at an individual or group who may have been the source of a particular exploit. It is far better to work together as a team, taking an enterprise-wide approach that may uncover problems that span across the enterprise. If the underlying cause of a structural problem allows holes in the dike to form, fixing one hole, without figuring out that weak cement throughout the dike is causing the problem, will only mean that the problem will pop up somewhere else in the future.

One method used in helping to determine root cause in an enterprise is using a technique that is found within the Six Sigma DMAIC (Define, Measure, Analyze, Improve, Control) methodology. This technique is known as the *5 whys* because you ask five why questions. Asking a why question, such as, "Why did our

database fail?” and then arriving at an answer, and then asking the next why question like, “Why wasn’t the new database version tested before launch?” eventually helps arrive at a root cause like peeling the layers of an onion. Five is considered a minimum number of questions, but many more can be added as needed.

AI analyzes solutions by determining a host of factors that not only includes expected results but also the cost of achieving those results.

A Smarter Adversary

Weaponized AI can be used by bad actors and nation states to penetrate computers and networks. There is no direct evidence that AI has been applied for nefarious purposes. But it’s such a possibility that most professionals consider it only a matter of time. Network intrusion and malware will be able to adapt in real time making detection nearly impossible and rendering signature-based detection moot.

The Internet, with all its billions of nodes, is an AI dream when it comes to analyzing the best places ripe for attack. Software probing for vulnerabilities coupled with smart phishing and improved social engineering attacks provide amazingly sophisticated attack vectors. Also, cybersecurity intelligence is primarily open, meaning that weaponized, cognitively intelligent, machines can analyze the same data, only with a different end purpose. Entire countries have applied massive resources to penetrating our private networks with the goal of impacting public opinion, thereby affecting policy decisions, and the ability to shut down our nation’s infrastructure on demand.

If your eyes have glazed over thinking (hoping) this has not become a reality, all you need to do is see what commercial advertisers have done using publicly available social media and web surfing information to target ads specific to individuals to completely understand what that power in the wrong hands can do. Computers can write articles as well as human authors can (this book was written by a real person, honest), constructing stories so believable that even well-informed people are fooled. Scared yet?

IN THIS CHAPTER

- » Providing insight into possible threats
- » Embodying understand, reason, and learn
- » Curating threat intelligence
- » Using QRadar with Watson for cybersecurity research and response
- » Winning with threat intelligence

Chapter 5

Using the Cognitive Capabilities of Watson to Investigate Security Incidents

IBM's Watson computer, best known for its 2011 prowess at *Jeopardy!*, has since taken on increasingly larger challenges. One of those challenges was being able to assist doctors at finding treatments for cancer. In some ways, the application of Watson to find solutions to cybersecurity threats is the same thing. Cybersecurity fights the cancer that eats away at our network infrastructure. Identifying possible root causes, finding solutions, and providing insight to cybersecurity professionals is Watson's latest domain.

Taking Intelligent Action

There is far too much cybersecurity intelligence data for a single person, or even a team of people, to digest. New information is created and published at an ever-increasing rate. Statistics at IBM

point to the fact that cybersecurity analysts are able to keep up with only 8 percent of newly published information. Without machine assistance, the battle against network intruders will be lost.

Watson is able to analyze billions of data points gathered from network security analysis programs and correlate it against all known structured and unstructured articles, threat feeds, books, blog posts, and other sources that provide cybersecurity intelligence. What Watson provides is insight into possible threats, and it can do that up to 60 times faster than attempting the same response without it. The report includes a list of possible threats and their ranking based on the likelihood that this is the threat being encountered. The synthesis includes collaboration with human engineers to perform analyses of root causes, particularly those encountered by the organization being attacked, and the vast amount of structured signature data, and lastly, insights gained by analyzing and learning from the cybersecurity corpus of information. The end result is ten times more actionable data than the analyst would otherwise have at his fingertips.

Watson's ability to process these vast sums of data fights something that has been called "cyber blindness." Applying the Watson Discovery Service to a corpus of knowledge allows for the discovery of deep connections that a human may never see. To create this corpus, Watson was trained with over three million documents and billions of data points.

Understand, Reason, and Learn

Chapter 4 introduces the idea of the acronym URL: Understand, Reason, Learn. Watson embodies this concept when it takes in a corpus of knowledge, wherever it exists, as long as the content is available digitally. After a corpus has been "digested," Watson is able to provide insights by finding connections and causality a human may have missed. This is the reasoning that provides the important insights a security operations center (SOC) needs to quickly shut down an intrusion.

The learning part is the secret sauce. Watson continually learns. It forms its insights based on cybersecurity data and learning done with human interaction the same way a human learns by having a partner. Watson sees about 100,000 updates a week. This include four million elements an hour and fifteen thousand documents a day.

AUGMENTED INTELLIGENCE

Big Data analytics is the 18,000-year-old science of analyzing stored data beginning with scratchings on bone sticks in Africa and continuing to the present day. The company later known as IBM used business intelligence to assist the United States government in analyzing the greatest single corpus of knowledge of its time, the 1880 census, reducing the analysis time of the census from an expected ten years to only three months.

While the history of augmented intelligence is rich and interesting, what has changed is a machine's ability to parse natural language, giving it the ability to "apprehend the relationships of presented facts" in a faster and far broader body of knowledge.

When researching information about cybersecurity exploits, Watson uses its ability to understand, reason, and learn across a vast library of information. Providing this information to humans "amplifies" the human's intelligence. Every attempt to manipulate data, whether it was using the abacus or Watson, has been an attempt to expand the information processing capabilities of people. Put simply, arming cybersecurity engineers with actionable intelligence not only makes them smarter but also about 60 times faster.

Watson has an ever-growing cybersecurity corpus that gives it constant new insights that provide decision support and augmented intelligence to SOC analysts. According to a February 2017 IBM news release, by that time IBM had already digested a million cybersecurity documents. Using the URL concept, cognitive computing appears to be the way forward in the cybersecurity war.

Applying Watson and Qradar

The software commonly used to analyze security information and events in an enterprise is called Security Intelligence and Event Management (SIEM). This software is a combination of two goals, managing events (SEM) and managing information (SIM). SEM software receives real-time security alerts from software applications (both on network and in the cloud) and smart network equipment. Log data, security records, and other information generated by a security event is stored and analyzed by the SIM portion.

QRadar is IBM's security analytics platform. Teamed with Watson's cognitive computing capabilities, it creates an adaptive and constantly learning SIEM. This has led it to assist in solving real-world cybersecurity challenges. By providing cognitive analysis of events and informational data, the time to squash a security event can be reduced from hours to minutes.

Winning with Threat Intelligence

“We are not alone.” This *X-Files* quote is a great way to introduce threat intelligence. There is no reason to re-invent the wheel or feel like no one else is experiencing the same security challenges. The fact is that most organizations are experiencing exactly the same challenges. Security analysts rely on the regular reporting of vulnerabilities and how they were solved or at least what actions people are taking to limit or overcome them. There is power in the crowd.

Threat intelligence is the collected evidence about a cyberattack that includes the context, mechanisms of the attack, actionable advice, and other information necessary to augment the intelligence of the SOC analyst. It is a mix of both structured and unstructured data. The structured data usually includes information such as the name given to the hack, signature of the exploit, possible log data, and IP addresses associated with it. The unstructured data will include things like a description of the mechanism, analyst notes, shared comments, historical information, research articles, security blogs and bulletins, and response advice.

EMPLOYING X-FORCE

X-Force Exchange is an IBM cloud-based threat intelligence platform that allows you to consume threat intelligence data from peers, share your own, and take action on this information. The AI portion comes from an ability to use natural language processing to consume this data and apply it to investigations of security incidents. Full information on X-Force Exchange is available at ibm.com/security/xforce.

IN THIS CHAPTER

- » Defeating ransomware
- » Looking at smarter AI in cybersecurity
- » Winning against malware and non-malware
- » Gaining a better understanding of neural and capsule networks

Chapter 6

Ten Trends in Cybersecurity

While reading these trends, know that Watson has probably already read them, and who knows, perhaps even smiled. Cognitive computing and all that it means has brought about a new understanding of artificial intelligence (AI). Its use in the field of cybersecurity has taken cybersecurity to a new level.

Responding to Ransomware

Ransomware was once one of the fastest growing cybersecurity threats. This trend began to slow in mid-2017, but anxiety remains high. There is even something called Ransomware as a Service (RaaS) so malicious people no longer even need to have technical skills to launch an attack. Everyone is susceptible, and the number of ways to make yourself vulnerable to such an attack is growing, and attackers are growing smarter. You can check Chapter 1 for more details on ransomware.



WARNING

Use only licensed software from reputable vendors. This is difficult to police in a large organization and must definitely be seen as a high-risk potential for infection. Remember that people are the weak link. Make them less weak by educating them, whether it's your employees or your family.

Combining Application Development and Cybersecurity

If your company writes software, whether it's desktop software, mobile applications, web applications, or the programs that run inside the Internet of Things (IoT) devices, you need to make cybersecurity part of your development cycle. One of the dangers in today's world of writing code is that many organizations use an agile development method to save time. The goal to be first to market often means your prototype code, once it's bug-free, becomes your production code, and the testing phase, if there is one, is either post-launch and done by beta test users or is shuffled in between releases.

If you aren't set up to do your own security testing, services exist that do that for you. They can do code evaluation and penetration testing. It's often best to have a whole new set of eyes on your code. It's very easy for programmers to become code blind or feel pressured to produce a finished product.

Using Deep Learning to Detect DGA-Generated Domains

Domain Generation Algorithms (DGAs) create pseudorandom domain names (sdlkfusdlfl.com). When malware phones home (attempts connection to a remote network for the purpose of command and control), it uses pseudo-randomly generated domain names to try and remain anonymous. The reason they are pseudo-random is that they're generated by using a seed that's shared between the malware and the attacker network it's trying to contact. These DGA algorithms can produce tens of thousands of domain names. Trying to blacklist them all is pretty much a waste of time because eventually one will get through and connect.

The trend is to use deep learning to recognize malicious domains created by a DGA. The randomness of the domain names, meaningless groups of letters, makes them fairly easy to use as training models for a neural network. After seeing enough of these pseudo-random domains, neural networks have been pretty successful at spotting them. Counterattempts to defeat neural networks are being made all the time using a host of strategies. A smart DGA keeps changing to stay ahead of attempts to thwart it.

This is a back and forth battle that will go on for some time. Neural network designers are getting smarter, too. The code they create is performing better at finding even the most clever DGA-generated domain name.

Detecting Non-Malware Threats

Examples of non-malware attack vectors include Locky, CryptoWall, CryptXXX, CTBLocker, and PowerWare. This is by no means a complete list. These attacks use vulnerabilities found in web browsers, Microsoft Office applications, and even more dangerous, operating system programs like PowerShell and Windows Management Instrumentation.

Detecting non-malware threats involves an after-the-fact look at computer behavior. Training neural networks and using machine learning algorithms to monitor normal behavior, and working with a cybersecurity analyst to help train the AI about what is and what is not normal behavior will help create better detection methods. See Chapter 1 for a deeper look at non-malware intrusion.

Adaptive Honeypots and Honeytokens

Honeypots and honeytokens are traps for hackers. These are computers, passwords, and other fake information set up on a network to begin the process of collecting information about the attack and ultimately the attacker. By allowing an attacker to believe he has gained access to a vulnerable machine, information can be collected. But, this is an old trick, and few attackers fall prey to it anymore. Enter adaptive honeypots and honeytokens.

An adaptive honeypot is one that changes its behavior based on the attack, thus luring the attacker to reveal as much information as possible. How this works is that the adaptive honeypot starts throwing up defenses as though it were a protected computer. How the attacker reacts tells the analyst a great deal about the attacker's skill level and the tools they use when confronted with each new challenge. It's like watching the attacker play a game. The result is that an AI solution gets to learn the behavior so that it's recognizable in the future. For more information on honeypots and honeytokens, flip back to Chapter 1.

Gaining a Better Understanding of How Neural Networks Work

Cybersecurity will continue to increase the use of neural networks to detect malware and non-malware attacks. In the human brain, neurons (nerve cells) are interconnected using *dendrites*, which are little fingers that attach to multiple other neurons at their axon terminal. *Electrical signals*, also known as *action potentials*, are sent between neurons as a means of communication. Once these electrical signals are received, a decision is made as to whether the receiving neuron sends that signal on to other neurons. Think of this like having a meeting and having enough votes to form a quorum, and then having an up or down vote. A successful up vote (buildup of electrical charge within a neuron) gets sent on, where the down vote or vote that doesn't meet the required number of votes to form a quorum simply gets rejected. This is known as meeting a threshold, or in biological terms, an action potential. Once a threshold is reached, a neuron depolarizes and is reset.

The end result of signals being processed by many layers of neurons is that only those that have continued to meet the threshold have continued to the next layer until, ultimately, an action is taken. For more information on neural networks, see Chapter 3.

Employing Capsule Networks

Neural networks are generally used to process large amounts of information. It is so large that, until recently, it wasn't computationally possible. But a process known as *convolutional neural networks* was developed that revolutionized image recognition. It's a long explanation, but I'll give you the single sentence idea. Convolutional layers filter the data flowing between layers in a neural network in order to reduce the amount of data processed by each subsequent layer. If the network were used to analyze an image, it would first find edges, then shapes, and finally objects in the image. The problem is that convolutional networks don't really understand where those objects fit in relation to one another. For instance, if a picture was determined to have two arms and two legs, it may not understand that the legs both go on the bottom and arms on top.

A *capsule network* is a network where the neurons in each layer are divided into capsules that represent the different properties of the same entity. These correspond to entities you are examining, such as the legs and arms from the previous example. The properties may include things like orientation, positional relation to other entities, color, and so forth when examining an image. This is a great improvement over convolutional networks in identifying entities no matter what orientation the images may appear.

At a time when facial recognition and other biometrics are being considered to replace passwords and two-factor authentication, it becomes important that the software doing the recognizing gets it right. While capsule networks have been used to examine images, the idea of relationships and entity information will only make neural nets smarter about identifying other complex patterns, like cybersecurity attacks. It's been said that capsule networks may completely replace traditional neural networks in the future.

Deep Reinforcement Learning

Reinforcement learning is something you might have done when training your dog. You reinforced the good behaviors with a treat and punished wrong behavior with a rebuke. This same idea is applied to neural networks tasked with winning a game. The reward comes in the way of points or successes leading to ultimately winning the game.

The difference between reinforcement learning and deep reinforcement learning is the depth (how many hidden layers are within the network). It's been shown that complex problems are best solved by adding more layers to the neural network, making them deep learning networks.

Deep neural networks use a process called *Q-Learning*. *Q-Learning* involves the terms *state* and *action*, which are used to find an optimal solution. In an example where rats are running a maze attempting to get the cheese, each chamber of the maze would be a state, and the direction a rat takes between states is the action. The rat remembers all its attempts. This builds up a matrix of success possibilities and eventually allows the determination of the best route.

Deep reinforcement learning gives AI the ability to learn unattended when faced with novel challenges. For example, put a robot

in a building and tell it to find the warehouse. It will wander around until it finally finds the warehouse. This task performed enough times will ultimately let the robot find the shortest route to the warehouse. In cybersecurity, the future holds a wide variety of novel scenarios in which AI will have to learn on its own the best way to defeat an attack. Cognitive computing will present insight into the best responses an analyst might take.

Protecting the IoT

IoT devices represent a huge expansion of the number of network endpoints that must be considered when protecting an entire network. Anything connected to the Internet that isn't an actual computing device, such as your BluRay player and television, your security camera, and your digital assistant, is an IoT device and another endpoint on the network.

Protecting IoT devices is a layered approach. First, the IoT operating system should be hardened. This comes with good security testing during development. Second, the network to which the IoT connects must also be hardened. The last layer is always the human layer. For instance, many smart TVs now include a web browser. Malicious programs downloaded to your TV can easily spread to other devices connected to the same network.

Predicting the Future

Cognitive computing assists in the investigation of cyber-threats and identifies the root cause of an attack. China has gone one step further in its attempt to battle criminal enterprise. It has begun employing AI to attempt the prediction of crimes before they happen. Li Meng, Vice-Minister of Science, announced that by using AI and facial recognition to identify people, Chinese authorities can gather information on people and their activities. Using Big Data, the crime-fighting AI is creating a rating system to tag highly suspicious groups of people. Not even masks will fool the facial recognition system. Trying to fool the system will only make it smarter by “re-identifying” an individual. This is unlikely to happen in a country where in most places won't even allow traffic cameras to catch speeders. But, it certainly points out a potential use for machines growing ever more intelligent.

AI is the future of cybersecurity

Both AI and cybersecurity are broad fields. This book helps you understand the technologies that make up AI and how AI is applied to the battle of cybersecurity. Cybersecurity is one of the greatest challenges of our generation. It seeks to protect data, ideas, and processes that hackers attempt to steal to prey on businesses and exploit people around the world. Help is needed to make security professionals more efficient and augment their intelligence. This assistance comes from AI.

Inside...

- Machine learning and deep learning
- Apply deep learning to cybersecurity
- Use Watson to investigate security incidents
- Ten trends in cybersecurity



Go to **Dummies.com**[®]
for videos, step-by-step photos,
how-to articles, or to shop!

for
dummies[®]
A Wiley Brand



Also available
as an e-book

ISBN: 978-1-119-50825-0
Part #: 42013142USEN-00
Not for resale

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.