Relativity® Trace

3 Steps to Building an AI-Based Surveillance Strategy

How to use artificial intelligence to better identify risk, remove obvious junk, and reduce false positive alerts.

A Communication Surveillance e-Book

For more than 10 years, Relativity has been building artificial intelligence (AI) solutions to sift through communications in highly regulated and scrutinized industries. Our software is used by more than 70 Fortune 100 companies, from regulators like the Commodity Futures Trading Commission to government agencies like the US Department of Justice. Relativity Trace monitors all communications in near-real time and alerts compliance officers to suspicious content.

When building and training AI solutions, no resource is more important than a wealth of data, and **in 2019 alone Relativity analyzed more than two petabytes of data. That's the equivalent of reading the dictionary more than 650 million times.**

At Relativity, we know that there is not just one algorithm, behavioral analytics tool, or machine learning model that will solve your surveillance challenges. That's why our dedicated AI team has honed more than 15 highly configurable capabilities that, when used together, will greatly reduce false positives and pinpoint organizational risk for compliance teams struggling to keep up with growing communication volumes.

Introduction

For highly regulated financial organizations with significant compliance risks, the current regulatory environment places a heavy burden to detect and take proactive action on noncompliant behavior exhibited by employees.

If compliance failures do occur, there is massive risk of legal, reputational, and financial damage. And while the obligation to detect abuse is increasing, so is the volume of data that compliance teams need to grapple with.

In today's workforce, employees leverage multiple tools to communicate and get their work done. You have financial chat platforms like Bloomberg, Reuters, and Symphony; enterprise chat systems like Slack, Teams, and Skype; and file sharing and collaboration applications like OneDrive, Box, and SharePoint. And that's before we consider phone calls, text, email, and all the attachments and images sent in those platforms.

So, what's your compliance team to do?

You need to be able to sift through all the information coming your way, reduce false positives, and zero in on the riskiest communications as quickly as possible.

To efficiently and accurately accomplish this goal, you need an AI strategy.

This e-book outlines the three steps that are essential to a successful AI strategy and details the individual AI capabilities necessary to reduce noise and pinpoint market manipulation, insider trading, collusion, and other high-risk activities. It's meant to challenge the way you think about an advanced surveillance practice and give you the questions you should be asking your communication surveillance vendor.

Step 1: Turn down the volume

Most of the data you process is irrelevant. It's duplicative, it's "system created," it's junk. Getting rid of this content is the most important part of understanding your data. If you're only removing exact duplicate files, you might as well throw all downstream AI out the window, as it will struggle to accurately identify patterns in the chaos.

There are multiple forms of deduplication that must be applied to your data to reduce duplicative alerts and increase reviewer speed. The most basic deduplications are exact-duplicate and near-duplicate identification. While you should always remove exact duplicates, near-duplicates should only be removed automatically if you have a small review team and are struggling to keep up with alert volumes. If this isn't the case and you plan to review near-duplicate documents, you should always serve them up together for reviewers and give them document compare tools so they can quickly identify the unique content and take action.

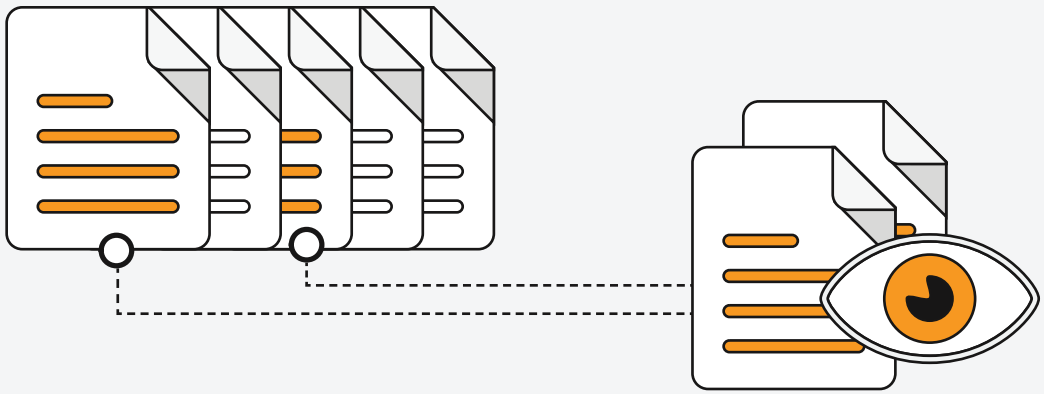

Although these basic forms of deduplication are important, email threading is the AI deduplication capability that has the greatest potential to **reduce your document counts and alert volumes by up to 60 percent.** Email threading is a text analytics feature that works behind the scenes to detect all emails in a single conversation thread. When emails are sent back and forth, each email in the conversation includes all the content from the previous emails in the thread. This means that if your surveillance tool alerts on content in the first email of a 20-email chain, that same content is going to trigger an alert in the following 19 emails. Think of the sheer number of alerts you would experience.

By applying email threading, not only will all emails in a single conversation thread be grouped and visualized together, but a single email that includes all the unique content from a conversation will be alerted on and served up to a reviewer. This drastically decreases the number of alerts, preventing repetitive work while ensuring all the content is reviewed.

When it comes to email threading tools, it's important to understand that not all are created equal. Email threading is extremely challenging due to branches in conversations and changing communication mediums. There are forwards, replies with added participants, attachments, and inline replies. A participant can also send emails using a desktop client, mobile device, or web app with different time zones, languages, and email addresses. We suggest that if you decide to add email threading to your AI strategy—which you should—test the capabilities before deciding on a provider.

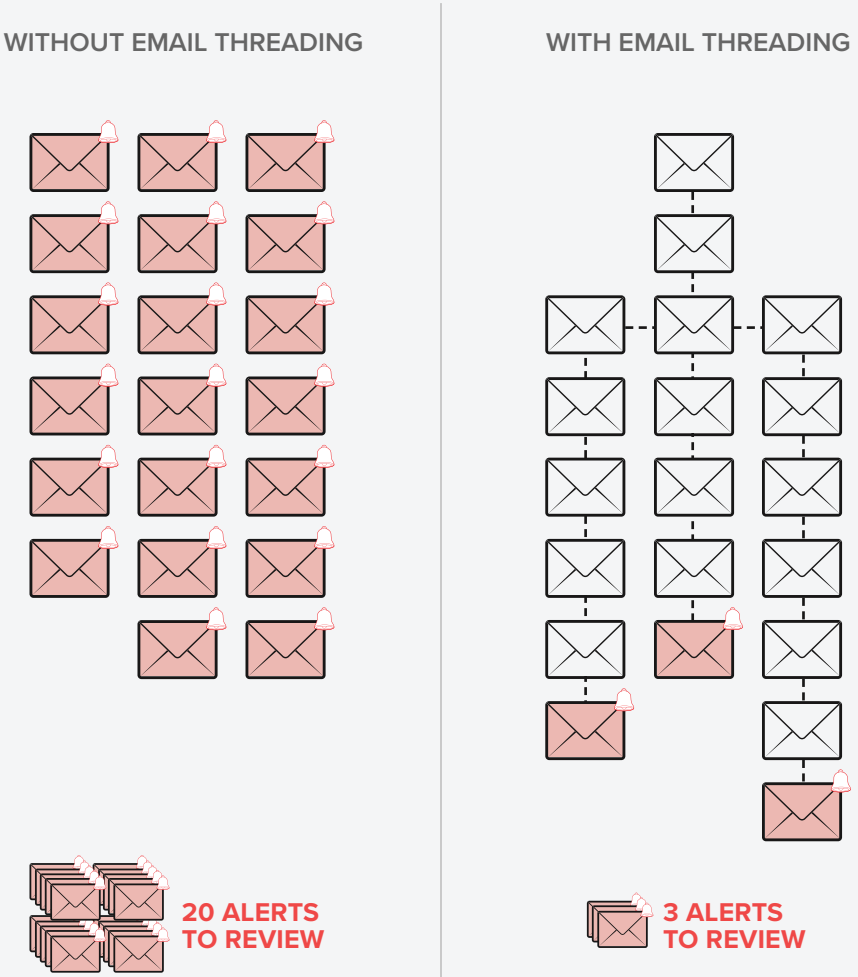

By applying email threading you can drastically decrease the number of alerts to review preventing repetitive work while ensuring all the content is viewed.

On top of deduplication, you need to eliminate system-created content from your emails. Disclaimers, email signatures, server stamps, confidentiality footers, and other repeated content will increase false positive alerting. References to different types of misconduct in confidentiality footers are notorious for causing large document volumes to be wrongfully alerted. Some tools attempt to strip this content out for you automatically, but this black-box approach can inadvertently remove relevant content. Rather, you want a tool that surfaces this content and gives you full control over what is removed and what is not. This ensures you have full transparency into the content that gets analyzed and allows you to spot fishy scenarios, like individuals hiding important content in confidentiality footers.

Spam disposal is the final approach to quickly removing irrelevant content. Spam in the corporate world is often different from what hits your personal account—it's made up of internal newsletters, blog posts, marketing content, and other mass messaging. To remove spam content, you need a pretrained model that can rank each communication based on the likelihood of it being considered spam. This ranking should allow you to set your own threshold for what qualifies as spam and should be removed. Relativity Trace not only ships with a tool to rank spam but also learns from newly identified spam to increase accuracy specifically for your organization. The more you use the tool, the more effective it becomes.

Reducing the noise in your data set is an essential first step in your AI strategy. The tools mentioned here and found in Relativity Trace will accurately and defensibly cut your compliance team's data set significantly.

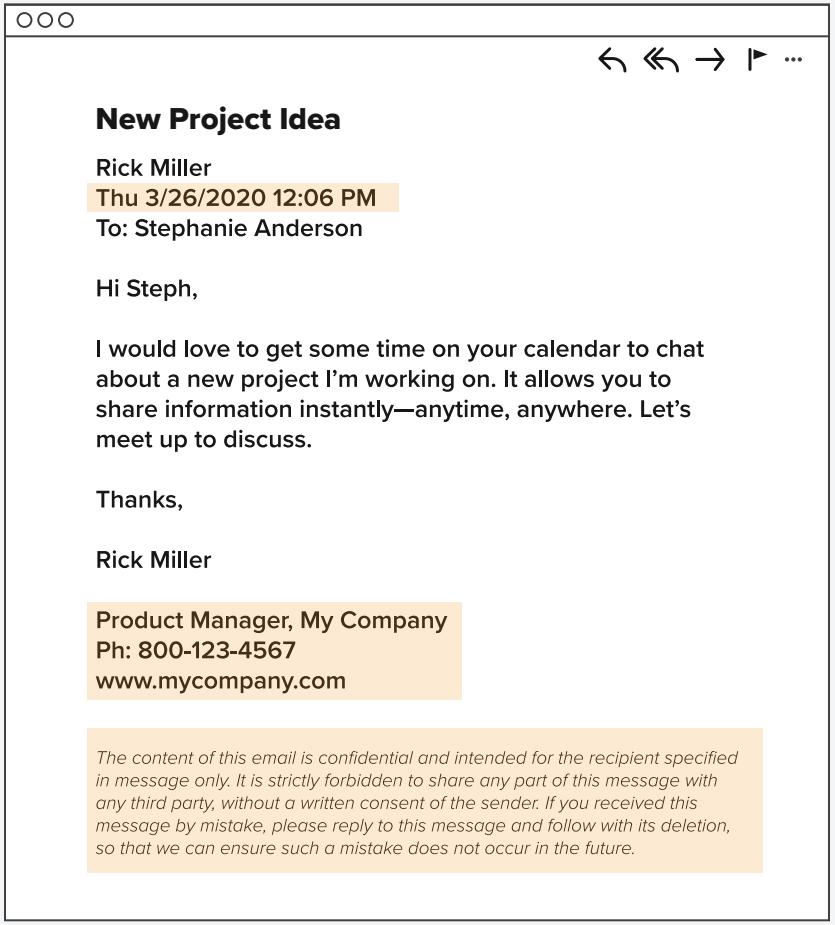

Quickly sift out system-created content such as disclaimers, email signatures, server stamps, confidentiality footers, and other repeated content to reduce false positive alerts.

Step 2: Find the signal in the noise

Now that you're focused on the communications that contain the most meaningful content, your system can start identifying and alerting on risk. At this point, many organizations like to throw a small list of terms (e.g. bid up, illegal, insider, avoid, corner, etc.) at their data and hope something sticks. Generally it does stick, but to communications that don't display any sort of risk (a.k.a. false positives). This rudimentary approach might meet some communication surveillance regulatory requirements in the short term, but it's not going to protect you from fines and reputational damage when misconduct is exposed years from now.

To alert on the riskiest content, you must lean heavily on unstructured and structured AI capabilities that identify patterns in your data, learn from past experiences, and make statistically appropriate decisions.

Unstructured AI capabilities can allow you to alert on the most complex file type, such as audio and images. You want an audio transcription tool that can be trained to understand domain-specific conversations in many languages and dialects. This is imperative when monitoring trader conversations where slang and abbreviations are frequently used. Relativity Trace is integrated with several audio transcription tools that allow for side-by-side audio listening capability and transcription tracking for easy review. For images, utilizing advanced optical character recognition technology to extract text that appears within the image can help you alert on screenshots that are commonly shared in chat and email data. Shady dealings are often pushed to mediums where monitoring is most difficult. Advanced AI in products like Relativity Trace make it so that there is nowhere to hide.

Unstructured AI capabilities can allow you to alert on the most complex file type, such as audio and images.

If you're a global business, you will need to identify the different languages used in each communication. This will not only allow you to direct specific documents to native-speaking reviewers, but also alert on monitored individuals who are trying to elude detection by switching between languages. If you don't have native-speaking reviewers, your tool should allow you to translate communications in real time. Whatever you do, don't copy and paste your confidential communications into Google Translate—make sure you use an in-platform translation tool that's secure.

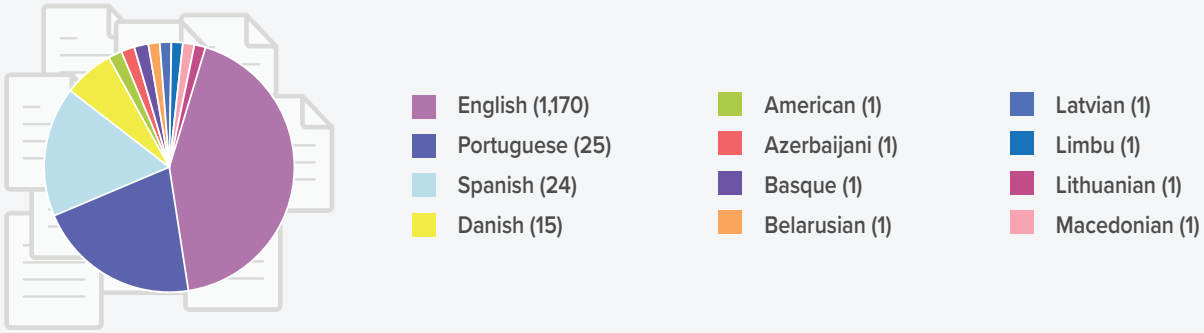

Other unstructured data analysis tools can provide insights from your data that you didn't even know existed. For example, if traders are using code words to hide discussions related to market manipulation or collusion, you would think it's nearly impossible to identify those communications without knowing the code words. That's where unstructured categorization AI capabilities come in. Tools, like Relativity Trace's conceptual clustering, categorize documents into named multi-depth groups based on the topics discussed in the text. This helps you identify topics that seem out of place—and it could be that code words are being used.

Step 2: Find the signal in the noise

Leveraging supervised machine learning capabilities that analyze unstructured data can provide instant risk ranking on new communications based on reviewer decisions from past alerts. As your compliance team flags documents as risky or not, the solution will continuously refine its understanding of what constitutes a compliance violation for your business.

Relativity Trace leverages fully integrated machine learning to score every new communication from 0 to 100 based on how risky the communication is, and your team designates what is a true or false positive. The risk scores are updated based on your team's review, and this process continues over time, making the machine smarter with every review.

Although unstructured data analysis can generate alerts on its own, pairing these results with structured data and advanced phrase searching will more effectively target risk and reduce false positive alerting. For example, you will be able to identify conversations in any language between your financial advisors and clients where guarantees are made regarding financial gain. You will be able to determine when market manipulation conversations are happening in the hours leading up to the market close based on past market manipulation scenarios you've witnessed. You will be able to catch when members of your investment bank are sharing non-public information with outsiders using code words that reference hockey players.

By combining your AI capabilities, extensive metadata filtering, and advanced phrase searching, you can alert on extremely specific situations where risk is most prevalent.

Step 3: Unravel an alert

For most compliance teams, an alert on a truly concerning communication is the catalyst that generates a case, kicks off an investigation, and requires communications with your superior. Your activities shift from review to an investigation to understand the context, intent, pressure, and opportunity around this event. A proper exploration must expand outside this single alerted document to similar communications and those between implicated parties. You may be required to pass your findings off to your legal department, where they'll collect data from custodians outside of your monitored individual list. Advanced AI investigation capability and visualizations are essential to quickly unearth the truth and take the appropriate mitigating action.

If the alerted document is an email, you will want to start by investigating the rest of the email conversation. Although the email you're reviewing will contain all previous content from that thread branch, there may be branches of the thread (such as forwards or replies) that don't appear in this email. There could also be inline comments or attachments in other emails in the thread group containing unique content that will be important when understanding context. Relativity Trace displays email thread information in a simple visualization that outlines how the conversation unfolded, highlights key variations in documents, and allows for reviewers to quickly jump from email to email.

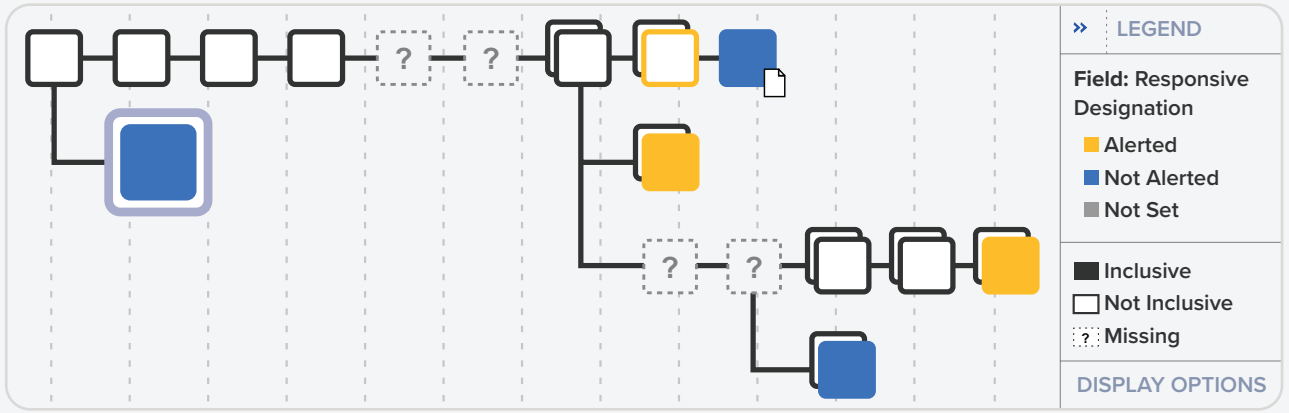

Relativity Trace displays email thread information in a simple visualization that outlines how the conversation unfolded, highlights key variations in documents, and allows for reviewers to quickly jump from email to email.

You should also look to documents that are conceptually similar in nature. These are communications that might have completely different text, but the topics being discussed are the same. Conceptual similarity capabilities will locate documents contextually like the complete text or specific sentence/paragraphs in the alerted document. This allows you to identify important documents even without knowing the specific terms, phrases, jargon, or code words that may be used by employees in your organization. Because potential misconduct is generally hidden in larger conversations, being able to examine specific sentences is extremely valuable when trying to locate similar content.

Conceptually similar documents

In short, conceptual analytics can find communications you're not even sure exist—a far more forgiving workflow than the typical constraints imposed by standard terms-based rules.

Analyzing the individuals participating in an alerted communication can lead you to other documents important to the event. AI that groups aliases (e.g. email address, phone number, etc.) to a single individual, paired with visualizations that map participants, can help you locate communications between two specific individuals or during off hours and can help you understand the participants' social dynamics.

A small number of compliance team explorations may require a deeper investigation by a legal team involving more custodians' data and a larger team of reviewers. Transferring not only the new data but also previous findings can be challenging if you aren't using a tool that supports both your compliance and legal teams. With Trace, built on the industry-leading Relativity platform, moving your case to a legal team can be done with the click of a button.

Use conceptual analytics to verify that the way your team has tagged a given document is consistent with how conceptually similar documents were tagged.

Conclusion

Increasing data volumes, dispersed communication platforms, and bad actors devising strategies to evade surveillance are making risk identification within organizations more challenging, and there is more pressure than ever on compliance teams to proactively catch misconduct.

With the proper AI strategy and a platform that offers an extensive suite of AI capabilities, individual surveillance team members are better equipped to remove irrelevant content, get alerted on the riskiest behavior, and unearth the truth.

The AI strategy discussed in this e-book is completely automated in Relativity Trace, so you can focus less on technical AI details and more on ensuring your organization eliminates risk exposure.

Take the next step. Let us help you enhance your communication monitoring.

trace@relativity.com or **+1 312.263.1177**

LEARN MORE

Relativity®