Portfolio Media. Inc. | 111 West 19th Street, 5th Floor | New York, NY 10011 | www.law360.com
Phone: +1 646 783 7100 | Fax: +1 646 783 7161 | customerservice@law360.com

Getting AI Tools Litigation-Ready Is Crucial For Finance Cos.

By **Alex Lakatos, Eric Evans and Reginald Goeke** (July 9, 2019, 2:50 PM EDT)

Alex Lakatos

Eric Evans

Reginald Goeke

If your artificial intelligence tools are not litigation-ready, then discovery in a lawsuit contesting decisions those tools have made could quickly become a nightmare: Your company may suffer enormous distractions and decreased productivity as it struggles to address litigation requirements that are inconsistent with its AI systems, data and culture; may be subjected to onerous court orders that interfere with its ability to conduct its core businesses; may even suffer adverse judgments on claims that lack merit.

Preparing for disputes is a crucial consideration for any financial services company using AI and big data to make important decisions, such as whether to extend credit to potential borrowers or whether to flag a transaction as posing an anti-money laundering or fraud risk. As every decision-maker knows, once you start making decisions, you cannot please all the people all of time. Thus, litigation over AI decision-making is not a question of if, but when.

This article provides a road map for addressing discovery challenges intrinsic to AI, long before any lawsuits are filed, early enough that a thoughtful strategy and modest investment of resources can have a butterfly effect, multiplying to enormous value when disputes later arise.

Below, we first discuss why AI poses unique discovery challenges, different in quantity and quality for those arising from prior disputes over computerized models and decisions. Second, we discuss how savvy plaintiffs lawyers will seek to exploit those challenges to obtain strategic advantages in litigation — particularly in today's world, where some rules governing discovery and evidence still lag behind technological realities. Third, we set forth practical, actionable steps that financial services companies deploying AI can implement now, to help mitigate serious problems down the road.

We note at the outset that AI is a rapidly developing field and that most litigation over AI has yet to occur. Moreover, judicial efforts to grapple with the unique challenges AI poses are nascent, or even nonexistent. No doubt the future holds surprises. Our experience in other, related litigation contexts informs the article throughout; however, financial services companies will be best served by taking a flexible, nimble approach toward applying the recommendations below.

Why AI Poses Unique Discovery Challenges

Machine learning systems pose potential discovery difficulties beyond those typical for conventional algorithms or computer programs because AI systems are different in three critical respects: (1) inputs, (2) processing and (3) outputs.

Inputs

Andrew Ng, former chief data scientist at Baidu and an often cited AI expert, analogizes deep learning models to rocket engines that requires loads of fuel that is data. Machine learning studies and learns from data: It is "trained" on data. That thirst for data leads to several discovery challenges.

Data Volume

"Most applications of artificial intelligence require huge volumes of data in order to learn and make intelligent decisions."[1] Moreover, as algorithms become more sophisticated, they require even greater amounts of data. If a linear algorithm — a comparatively simple approach to machine learning — "achieves good performance with hundreds of examples per class, a nonlinear algorithm may need thousands of examples per class."[2]

Indeed, AI often functions by analyzing all the data that is available, e.g., reviewing all transactions, customer data, behavioral data and the like to spot money laundering risks or to assess creditworthiness. Producing and reviewing this data, as litigation often requires, poses significant challenges.

Data Sensitivity

In many instances involving financial services companies, the data used to train the AI will be sensitive. The data may include personally identifiable information, such as social security numbers and date of birth. It may reveal an individual's financial health and personal spending habits. It may contain medical information, such as spending on health professionals. In some cases, the financial institutions may owe duties of confidentiality to their customers. In other cases, while no official obligation may exist (or while obligations may be subordinate to discovery production obligations), the financial institution may still wish to protect its customers' privacy, whether for reputational reasons or as a matter of its own corporate values. This too, creates a challenge, especially when, as discussed above, vast troves of data are at issue.

Data Evolution

Machine learning systems may be designed to learn iteratively, refining their decision-making every time they receive additional data. A machine learning system that recommended extending credit on day one might make a different recommendation on day two, based on the system having seen more data, and having learned more, and having refined its internal model, in the interim. This presents discovery challenges for data. For example, is it even possible to go back and identify the data that the machine learning system trained on at a particular moment?

Data Retention

Many AI systems overwrite training data to conserve storage and other resources. Given the vast

volume of data, and the fact data often ages out of usefulness, it may be impractical to maintain the data that led to a decision. But litigation-related preservation obligations do not automatically take practicality into account. Determining which data may be overwritten, when and how it might be preserved, and ensuring that the space exists to preserve it, can pose a significant challenge.

Processing

The manner in which AI tools analyze data to reach decisions is, more than any other factor, what separates AI from prior, algorithmic decision making programs. Those differences, however, create a host of discovery challenges.

The "Black Box"

Before modern machine learning, algorithms made decisions that were rules-based, and so could be understood by studying the rules (the computer code) behind them. For example, a rules-based program might have a rule that provides that if borrower has a certain debt-to-income ratio above a certain amount, then the lender will not extend any additional credit to that particular borrower. Machine learning lacks such deterministic rules. Machine learning is probabilistic and uses statistical models. Machine learning might approach that problem above by building a model to answer the question: how much does this potential borrower resemble those who have paid-as-agreed, versus those who have defaulted?

The more complex the machine learning algorithm is, the more opaque the model and the harder it is to know why it made the decisions it made — e.g., what factors it weighed, how much weight it gave those factors, and how those factors interrelated. For example, machine learning easily beats human grand master chess champions, making moves that would not occur to them, based on reasoning they cannot fathom. Machine learning tools don't "think" like people: They have an "often quirky imagination."[3]

Output from a machine learning system may offer only limited insight into what is happening inside the black box. Machine learning is only concerned with the specific outcomes that its engineers instruct it to care about. It therefore may take a path to get to an end point that humans would consider to be cheating, undesirable or otherwise inconsistent with their intentions.

Sometimes, the outputs may instantly reveal that the machine has taken an unanticipated path, such as the following real world examples: (1) a robotic arm trained to slide a block to a target position on a table instead achieved the goal by moving the table itself and (2) an artificial life simulation where survival required energy but giving birth had no energy cost, one species evolved a sedentary lifestyle that consisted mostly of mating to produce new children to devour.[4]

On the hand, sometimes it will be far less apparent from the output that something unanticipated is happening inside the black box, such as a case in which AI trained to classify skin lesions as potentially cancerous learned that lesions photographed next to a ruler were more likely to be malignant.[5]

Data Retention Within the Black Box

The probabilistic decision-making model that comprises an AI tool (the inner workings of the "black box") may change from time-to-time, or even iteratively, as the AI learns from new data and adjusts the model accordingly. The AI system, however, may not be configured to retain values that change or are overwritten as the AI learns. For example, for deep neural networks, the value of the weights in each

node may not be something the system has any means to preserve as the system repeatedly readjusts and refines them.[6]

AI Development

AI systems that are of significance to an entire organization, if not mission critical, typically are customized by and for the organization, and may be fine-tuned to reflect specific data sources, customers, marketing strategies, products and the like. Systems may become so complex that even their creators have trouble understanding them. Such systems may have rules of operations different from products offered on the mass market.

Moreover, such systems are often in a process of constant updating and revisions by software engineers and data scientists, who may be doing anything from experimenting with new techniques for analysis to tweaking the inputs or outputs. There may not be one static set of code to produce, but millions of lines, with hundreds or thousands of owners, in a constant state of flux.

In addition, older machine learning systems may be superseded by newer versions. Once a system is no longer in active use, it may be difficult to maintain it in a usable form. On the other side of the same coin, many new lines of business do not succeed. If anything, this may be more true when it comes to new AI ventures, as the technology is novel and complex, and the regulatory environment is unsure.

The challenges described above are exacerbated when an AI system is shelved, or heading toward the dustbin. At that point, the incentives to preserve the system in working order, and to maintain data and records about the system, are diminished. Yet just as incentives to maintain and preserve the system are waning, the risk of litigation about the system may be waxing.

The more heavily documented a complex system is (with changes, modifications and even basic functionality documented and explained in real time), the more feasible it will be at a later date to explain the system and have a production (e.g., of code and training data) that is meaningful. But computer scientists and software engineers, particular those in nimble fintech startups, may not have a strong culture of documentation, and may even consider documentation inconsistent with their “flash on insight” programming methodology.

And even if computer scientists and software engineers routinely document their AI systems, that documentation may be more geared toward the needs of other computer scientists, not of litigants who need to explain AI systems for litigation purposes. In a high paced development environment, and one where turnover of AI programmers (who are in high demand) is routine, it is not uncommon for system documentation to consistent of shorthand and over simplifications.[7] But in a discovery context, plaintiffs’ counsel may take such shorthand or simplifications out of context, argue that programmers are trying hide their decisions, or worse.

The "Secret Sauce"

Institutions that utilize AI often consider the exact programming and training of the AI an important trade secret. Some organizations, especially vendors supplying machine learning systems to financial institutions, may even promote their AI as better than those of others for reasons they cannot share (their “secret sauce”). Balancing the desire to protect valuable intellectual property from plaintiffs lawyers’ who are likely to demand maximum discovery — both to help prove their case and, sometimes to help coerce a settlement — is yet another discovery challenge.

Outputs

Size and Complexity

Some AI systems have outputs that are huge and complex, and that can take a long time to generate. An AI tool that looks for fraud, for example, may review and risk weight thousands of transactions per hour. Further, the output may not be user-friendly; it may require expertise in a particular system to understand the meaning of the AI output. And that meaning may change over time, as the metrics for scoring or the interface are amended to reflect ongoing developments.

The Panda Problem

Pandas are animal known for living comfortably in their native habitat, but doing poorly when transplanted to another environment. Similarly, AI outputs may be usable and comprehensible within the system surrounding the AI tool, but may be hard to export in a meaningful fashion outside that native software environment.

For example, the information that AI system outputs may be (1) stored in deep storage, so that it first must be moved to fast storage before it can be searched, collated and utilized, (2) stored in a proprietary format that is exotic, as opposed to commonly known formats, such as .xls or csv; (3) subject to search and review only using specialized tools that may only exist in-house, and that may be understood only by in-house engineers.

Which of the problems above presents the biggest risk for your AI system?

Depending on how you deploy your AI, some of the discovery risks above may present greater challenges than others.

Predictive Analytics Tools

In the case of predictive analytics tools, such as AI that performs credit scoring, dispute are likely to focus on how the tool was trained and how it made decisions — which implicates the "black box" issues discussed above. For example, how does the black box predict credit performance? To what extent does it rely upon impermissible information, such as factors that are closely correlated with race, gender or other protected characteristics? And to what extent did the training data include such proxy information? Another contentious area for this type of tool may be quality data, and the possibility that inaccurate data led to inaccurate scoring — which implicates the data input problems discussed above.

RegTech AI Tools

In the case of regulatory technology AI tools, such as those used to mine the company's own data to identify trends or regulatory violations (e.g., Bank Secrecy Act violations, patterns in customer complaints), some disputes may focus on what the financial services company knew, and when the financial services company knew it — which implicates the data output challenges described above.

Other disputes may involve plaintiffs attorneys who want to mine the data themselves, and look for new problems as fodder for new or expanded claims. To do that, they will want all of the data input and their own access to the tool— which implicates the data input and black box problems. The same is true of disputes over whether the financial institution could, and should, have spotted a problem that it failed

to identify, such as a Ponzi scheme orchestrated by one of the financial institution's customers that injured civil plaintiffs suing the financial institution.

Portfolio Development AI Tools

In the case of portfolio development AI tools, designed to improve the performance of a pool of assets (e.g., a loan portfolio, a hedge fund), disputes are likely to arise in cases of underperformance and focus on allegations that owner/developer of the AI misrepresented its features and capabilities. A party that wishes to demonstrate that the tool did not perform as advertised (and that fault lies with the tool developer) may seek to scrutinize the tools outputs, its training, development and inner workings — which implicates black box and data output issues.

An adverse party that wished to demonstrate that the tool worked well, and underperformance is due to misuse (and that fault lies with the tool operator) may seek the scrutinize the parameters that the tool user adjusted and the quality of the data that the tool user input — which implicates a different set of concerns more focused on inputs.

Marketing/Sales AI Tools

These tools, which may mine customer data to enhance customer service or to identify marketing and sales targets for particular product, or which may interact directly with customers (e.g., chatbots), may raise yet another set of challenges, e.g., data input related in the case of data mining.

How Plaintiffs Lawyers Will Seek to Exploit Discovery Challenges for Strategic Litigation Advantages

Below, we discuss several areas where the perfect storm of discovery jeopardy may arise from the intersection of (1) complex modern AI, (2) rules of civil procedure (and common law guidance) that in some respects may lag technological developments and (3) aggressive plaintiffs lawyers. In particular, we discuss issues relating to preservation, production, and proof.

Preservation of Documents

The Risk: Sanctions

"Spoliation is the destruction or significant alteration of evidence, or the failure to preserve property for another's use as evidence in pending or reasonably foreseeable litigation."[8] Generally, trial courts have broad discretion to impose sanctions for spoliation. Depending on the forum, sanctions may include, among other things, precluding a party from presenting evidence on topics addressed in evidence that was subject to spoliation, allowing evidentiary inferences that the missing evidence would have been adverse to the party that failed to preserve it, finding certain issues conclusively established against the party that failed to preserve evidence, and entering a default judgment against the party responsible the spoliation.[9] In determining what sanctions (if any) are appropriate, courts generally will consider whether the party's preservation efforts were reasonable and undertaken in good faith.

In 2015, in recognition of the complexities and difficulties of preserving electronically stored information and the overzealous application of sanctions in certain cases, the Federal Rules of Civil Procedure were updated to provide guardrails constraining the ability of federal courts to impose sanctions for spoliation of ESI.

Specifically, under Rule 37(e)(2), a court must conclude that a party's spoliation of ESI was intentional (and not merely negligent or grossly negligent) before imposing more serious discovery sanctions, such as an adverse inference or default judgment. This rule helps protect parties utilizing AI tools, although some courts have been willing to infer intent from the circumstances of the preservation failure itself,[10] and some courts have allowed the jury to decide in the first instance whether the failure to preserve ESI was intentional.[11] State courts also may allow discovery sanctions for the spoliation of ESI on a showing of negligence or recklessness.[12]

How Plaintiffs Attorneys Will Try to Exploit the Situation

As detailed above, preserving incoming data streams, data explaining the working of AI and data output, all pose challenges. Plaintiffs are likely to try to exploit these difficulties in various ways, including:

- Sending letters making unreasonable demands for document preservation at the outset of the litigation, and revisiting the issue during meet-and-confer discussions on electronic discovery;

- Seeking discovery of documents (such as manuals, code and code documentation) and taking depositions (e.g., of corporate representatives, of engineers) to test what preservation might be feasible;

- Employing "experts" who will take unreasonable and unrealistic positions about what preservation is possible;

- Exploiting a court's lack of understanding about AI, playing on the common misperception that preservation of computer data is largely a straightforward matter (e.g., precluding automatic overwriting of certain data), and by arguing that preservation is simple and inexpensive to achieve; and

- Asking the court to infer that any data not preserved was a product of bad faith and seeking discovery sanctions.

Production of Documents

The Risk: Sanctions

Generally, parties must produce relevant information as part of the discovery process, typically in response to discovery demands. The production obligation is not unlimited, however. In federal court, for example, documents requests must be proportional to the needs of the case.[13] Similarly, state courts generally permit objections to discovery requests that are unduly burdensome. In disputes over production of AI, however, there is little guidance over where the proportionality/unduly burdensome line should be drawn.

To the extent the parties cannot agree on what data about an AI system will be produced, the next step will be seek guidance from the court. Once the court has determined and entered an order governing the required scope of discovery, failure to comply may lead to discovery sanctions such as those described above.[14]

How Plaintiffs Attorneys Will Try to Exploit the Situation

As detailed above, producing training data, decision-making data, data describing the inner workings of an AI system, and AI system outputs all pose many challenges. Plaintiffs are likely to try to exploit these difficulties in various ways, including:

- Making unduly broad discovery demands;
- Making discovery demands the seek sensitive customer information;
- Making discovery demands that seek trade secrets related to the functionality and operation of the financial services institution's AI tools;
- Arguing that the financial institution has possession, custody or control of information concerning AI tools that resides with the financial institution's third-party vendors, and that the vendors are reluctant or unwilling to provide to the financial institution;
- Arguing that any materials produced are inadequate;
- Requesting that their experts be afforded direct, onsite access to the financial service's AI systems;
- Seeking court orders requiring the financial institution to produce data that goes beyond what is practical, and perhaps even beyond what is possible; and
- Filing motions for sanctions.

Proof Based Upon AI-Generated Evidence

The Risk: Inability to Present Evidence to Support a Defense

Generally, to have an AI model, or output generated by an AI tool, admitted in evidence in a judicial proceeding, the party presenting the evidence must "authenticate" it; that is, the party must demonstrate the evidence is what it purports to be.[15] "Among the factors courts may apply in determining whether a proper foundation for admission of computer-generated evidence has been laid include whether the computer was standard and in good working order, whether the operators of the equipment were qualified, whether proper procedures were followed, whether reliable software was used, whether the program operated properly, and the exhibit derived from the computer."[16]

In his treatise on federal evidence, Judge Jack Weinstein explains the rigor required to authenticate computer-generated evidence will depend in several factors, including (1) the quality of the data input, (2) the complexity of the algorithm, (3) whether the problem is routine or novel, and (4) whether the output can be tested and verified.[17]

Recent amendments to the federal rules streamline the process for authenticating "a record generated by an electronic process or system that produces an accurate result." That rule is intended for routine computer-generated evidence, such as electronic phone log. By contrast, AI models with inputs, weights and outputs that are in flux, or that are novel and hard to comprehend, may encounter authentication challenges.

How Plaintiffs Attorneys Will Try to Exploit the Situation

As detailed above, authenticating AI models and outputs for admission into evidence may be challenging. Plaintiffs are likely to try to exploit these difficulties in various ways, including:

- Seeking discovery of all facts that may bear on authentication;
- Filing overreaching discovery requests on issues they maintain bear on authentication;
- Retaining experts to challenge the authenticity of evidence; and
- Filing motions in limine to exclude a financial institution's AI-related evidence.

Approaches to Mitigating Discovery and Evidentiary AI Risks

Explainability

The better you can explain your AI tools — what the tool considered and why, how the tool made its decision and why — the stronger your position will be in a litigation disputing its decision and the effect of its decision.

Explainability, in general terms, has three aspects:

- Transparency: easy identification of the important factors in the tool's operation;
- Interpretability: easy identification and explanation of how the tool weights those factors and derives them from its input data; and
- Provenance: easy identification of where input data originates and what the data contains.

We recommend a two part approach to defending the actions of AI tools. First, be prepared to frame the discussion in terms of decisions people made. In other words, AI users will want to prove that the business decisions are the policy choices made by company management and the choice of algorithms and parameters by data scientists and programmers based on those policy decisions. The AI is not a decision-maker but merely a mechanism for implementing those business decisions. To do this, ensure the following:

- The management team specifies to its data scientists, computer scientists, software engineers and technicians how the company wants to tool to work, recognizing that those specifications are in fact the business decisions;
- The company's e-discovery/information governance team specifies to data scientists, computer scientists, software engineers and technicians how the company wants to store and access input, outputs, change logs, models and the like;
- The company consults its legal and compliance teams on the points above;
- The company employs "AI sustainers" to continual test and modify the tool to keep it working as the management team and the e-discovery/information governance team intended; and
- The company employs "AI explainers" who know how explain the tool's results.

Second, based on the guidance described above, the company can include features in its AI tool that further support explainability, such as:

- Code that permits auditing and testing;
- Explainable AI; i.e. the cutting edge, and still nascent, techniques that are beginning to allow a window in the AI "black box";
- Extra documentation that explains how the AI works and what choices were made about its features and functionality, for the benefits of current in-house employees, later in-house employees and later retained experts; and
- Thoughtful decisions about what facts and data to preserve and which to overwrite.

Explainability will be invaluable when confronted with the problems of production, preservation and proof described above.

First, when it comes to preservation of documents and data, explainability affords the company several advantages: (1) the advance creation of relevant records, such as system documentation, (2) an understanding of which records are important, so that the AI tool and related policies and procedures can be designed to better preserve those records, and (3) positioning the company to defend its choices about what and how to fulfill its document preservation obligations.

Second, and relatedly, explainability aids with document production by helping to (1) ensure that key records are created and preserved, as described above, (2) prepare the company and that AI tool to export data in a comprehensible, portable format, (3) position the company to argue that its production is appropriate and to defend against overreaching or misguided demands for other information or access.

Third, explainability goes to the heart of authentication: the better the company can explain its tool, the better the company can demonstrate that the AI model and outputs are what they purport to be.

Looking beyond discovery, to the merits of the case — i.e., when the company must justify the decisions of its AI tool — explainability will once again inure to the company's benefit. AI systems that aren't designed for explainability are often difficult to defend. Plaintiffs will provide expert testimony stating that the AI should have resulted in a one set of decisions that, not surprisingly, establish that they were harmed. In response, defendants will proffer their own evidence, usually from an expert, that tries to show that the AI functioned as intended and plaintiffs were not harmed.

But this battle of experts is fundamentally biased toward plaintiffs. When their experts are not presented with an explainable AI, they can simply provide favorable assumptions to traverse any difficult-to-explain aspects of the AI. They can do this because, as a practical matter, once an AI becomes too complicated to explain elegantly, a finder of fact will default to the simpler and cleaner explanation of the messier and more complicated one. Further, teams that create AI tools that aren't optimized for explainability will often throw off statements about "fixing" a "broken" AI system that feed directly into plaintiffs' narratives and undercut defendants.

An AI system optimized for explainability, though, can become almost a witness in its own defense. Design documents written for an audience of regulators or end users will make defendants' points

better than design documents written for doctors of computer science. Data retention decisions can provide the key data points required to demonstrate the operation of the system instead of leaving it a black box. Documented design meetings and reports from sustainers can give the defendant human-scale stories, in human language, describing the AI — as opposed to mountains of raw data and near-indecipherable source code. And that change can level the playing field for companies defending their business-critical AI systems.

Storage of AI Tools and Information

It is also important to consider, at the outset of an AI project, where the tool will reside. There are several advantages to on-site storage from a litigation-ready perspective. First, litigation may require decades-long retention of data stores, which can add up. Second, if you own a server, you can always turn it off and physically shelve it, if necessary to preserve a legacy system.

As for data, consider whether storage of higher risk data is necessary to your project. For example, if there are categories of personally identifiable information, health information or financial information that you do not need, consider whether you can avoid collecting and keeping that information. In the alternative, consider whether it is possible to anonymize information so that it's no longer personally identifiable.

Conclusion

Because AI tools are becoming ubiquitous in the financial services ecosystem, and because AI tools are more involved in decision-making than their predecessors, you should anticipate a flood of AI-related disputes. Companies that fail to prepare may find themselves drowning. Those that ensure that their AI is litigation-ready, by contrast, are well positioned to stay afloat. Now, when the levies have yet to break, is the time to act.

Alex C. Lakatos, Eric B. Evans and Reginald R. Goeke are partners at Mayer Brown LLP.

[1] Artificial Intelligence and Privacy, Datatilsynet (Norwegian Data Protection Authority) at page 4 (January 2018), available at https://www.datatilsynet.no/globalassets/global/english/ai-and-privacy.pdf.

[2] https://machinelearningmastery.com/much-training-data-required-machine-learning/. Non-linear algorithms like a random forest or an artificial neural network are more sophisticated approaches to machine learning.

[3] https://www.technologyreview.com/s/612898/ai-is-reinventing-the-way-we-invent/

[4] https://boingboing.net/2018/11/12/local-optima-r-us.html

[5] https://www.wired.com/story/when-bots-teach-themselves-to-cheat/

[6] A node combines input from the data it receives with a set of weights, that either amplify or dampen

that input, thereby assigning significance to inputs with regard to the task the algorithm is trying to learn; e.g. which input is most helpful is classifying data without error? https://skymind.ai/wiki/neural-network.

[7] Nobel Prize winning physicist Richard Feymann tells a story from his time working on the atomic bomb in which junior physicists (many later recognized as geniuses of our time) were working so feverishly on a computer problem — at the time, based on punch cards getting out of order—that they didn't have time to explain the situation to their supervisor, much less document. His boss turned around and walked out of the room rather than interfere with the problem solvers working on the ground. Over 50 years later, his insight into how front line programmers see the world still rings true.

[8] West v. Goodyear Tire & Rubber Co.,167 F.3d 776, 779 (2d Cir. 1999).

[9] See e.g., FRCP 37; [state cites].

[10] See, e.g., O'Berry v. Turner, 2016 WL 1700403 (M.D. Ga. Apr. 27, 2016) (holding that failure to preserve ESI, reliance on a single hard copy, and loss of that hard copysupported a finding of intent to deprive).

[11] See, e.g., Cahill v. Dart, 2016 WL 7034139 (N.D. Ill. Dec. 2, 2016) (holding that the jury should make the decision whether prison officials had intentionally allowed a crucial party of a videotape segment to be overwritten).

[12] [cites].

[13] See FRCP 26(b)(1).

[14] See, e.g., FRCP 37(b)(2).

[15] See Fed. R. Evid. 901(a) ("The requirement of of authentication or identification as a condition precedent to admissibility is satisfied by evidence sufficient to support a finding that the matter in question is what its proponent claims.")

[16] https://www.americanbar.org/content/dam/aba/administrative/labor_law/meetings/2011/ac2011/123.authcheckdam.pdf

[17] Jack B. Weinstein & Margaret A. Berger, Weinstein's Federal Evidence at § 900.06[3].