Advantages of real-time quality control (and how to achieve it)

As Peter Drucker once said, "There is nothing so useless as doing efficiently that which should not be done at all." At DISCO, we've taken this sentiment to heart when it comes to quality control (QC). Rather than building multi-step post-hoc QC workflows, we've designed QC tools to help you avoid making mistakes in the first place.

Ask most legal professionals how they perform QC and they'll describe the frantic pre-production process for detecting errors in documents intended for production, or worse, the complex "recipes" and scripts needed to complete QC before documents go out the door. Because these steps are complex, time-consuming, and often incur vendor fees, even savvy legal professionals tend to save them until the end to avoid performing them multiple times.

But postponing QC until production time adds unnecessary stress to an already stressful process. Delaying QC also allows early errors to proliferate throughout a document set, creating more work when you are already crunched for time.

DISCO's real-time QC capabilities help you avoid both of these problems. By creating a constant feedback loop, you can perform efficient QC without waiting until the end of your review, and catch errors early to ensure there is less QC work to perform.

1. Identify review(er) errors as soon as they happen

The earlier you can identify mistakes or outliers in coding, the better.

With DISCO, you can identify anomalous reviewer behaviour (and possible misunderstandings of review protocol) by using built-in review metrics to monitor tagging behaviour. You can check for these trends at any time during a review, rather than waiting to pull a report at the end of the day or week.

For example, if one reviewer is tagging documents for Issue A significantly more often than other reviewers, DISCO will highlight the outlier. You can then review that reviewer's decisions to make sure they understand that issue. Similarly, if all reviewers are tagging documents with Issue C at a rate higher than you expected, you can create a target sample of those documents across all reviewers.

By using review metrics to look at reviewer and tagging behaviour, DISCO users can get quick insight into anomalies and provide feedback to reviewers faster. A tight feedback loop with the review team leads to better decision making in first pass review, reducing the work required during QC.
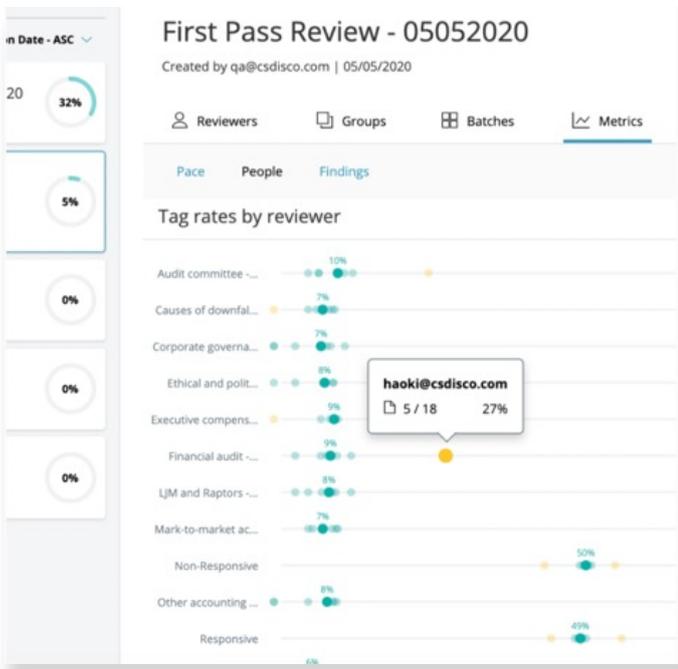
2. Automatically surface possible errors with DISCO AI

Deciding which documents to QC shouldn't feel like a guessing game. Not only are you potentially missing key documents, you're also spending more time to find the documents you do catch.

Instead, you can use DISCO AI to identify which documents are most likely miscoded. Using the same algorithm that allows it to prioritize responsive documents during review, DISCO AI can flag likely mistakes by identifying for which documents the reviewers' decisions and the AI model's prediction differ. In this way, DISCO AI can greatly reduce the time you spend identifying which documents need to be QCed.

All you have to do in this AI-driven QC model is open up your filters or search visualization to see if there are any coding discrepancies. With another click, documents are pulled up and instantly reviewable.
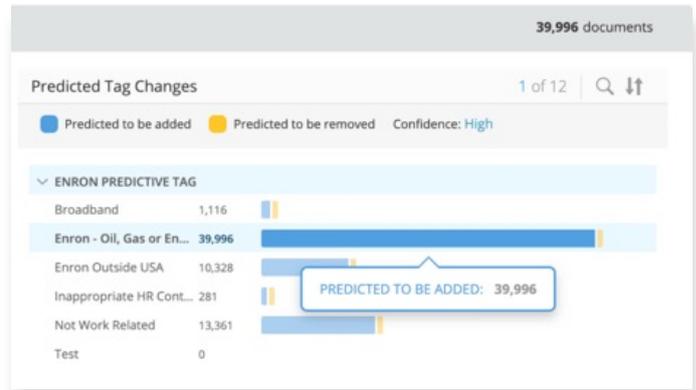
This simple QC workflow can be employed frequently even in the midst of a review, allowing the QC team to provide real-time coaching and feedback to reviewers.

Unlike the average document review where more than 20% of the total document population will go through some sort of QC, AI-driven QC results in more accurate coding during the review, which can reduce the QC population by as much as 80%. What is more, reviews using this methodology typically result in overturn rates of less than 3%, even when random sampling is done to double-check the AI's work.

3. Easily find inconsistently or consistently coded families

Whether you're working in an AI-prioritized workflow or a traditional linear review, inconsistently coded families can cause headaches. DISCO (and most ediscovery workflow experts) suggests coding documents on the face of the document alone — rather than at the family level — to allow for better AI learning. But this introduces challenges when it comes time to produce. How do you ensure that you pull out all documents that are privileged or confidential because of their family members before production?

In many platforms, this requires running multiple complex searches, creating folders, and then running exclusion searches. Depending on the size of your database, running these searches could take hours, if not days, to return the documents you need to QC.

With DISCO, you can identify these documents in seconds. DISCO has simple, elegant search syntax that pulls up inconsistently coded families with one click: familyInconsistentTag(_). Typing a tag like "privileged" between the parentheses and turning on family search pulls up all families that have some documents tagged as privileged and some without that tag.

This syntax creates the ability to easily have a secondary review for production running in parallel with your first pass review. By performing this review in tandem, you can eliminate the stress of trying to identify privileged or confidential documents on the eve of your production, which allows you to focus on more substantive QC or (gasp!) get home early and read your kids a bedtime story.