

How machine learning and content reviewers enforce policies

With more than 500 hours of video uploaded to YouTube every minute, it's important that we find and remove content that violates our policies at scale. To protect the YouTube community, we use a combination of advanced machine learning systems and expert human content reviewers to detect, review, and remove content that violates our policies before it's widely viewed.

What is machine learning?

Machine learning is a form of artificial intelligence whereby computers are trained to take actions without being explicitly programmed to do so. At YouTube, we rely on this technology to learn and identify patterns in content that may violate our Community Guidelines. These systems scan content on our platform 24/7, enabling us to review hundreds of thousands of videos in a fraction of the time it would take a person to do the same.

Proactively identify and flag harmful content at scale around the clock

Automatically remove content similar to what's been previously removed - such as spam

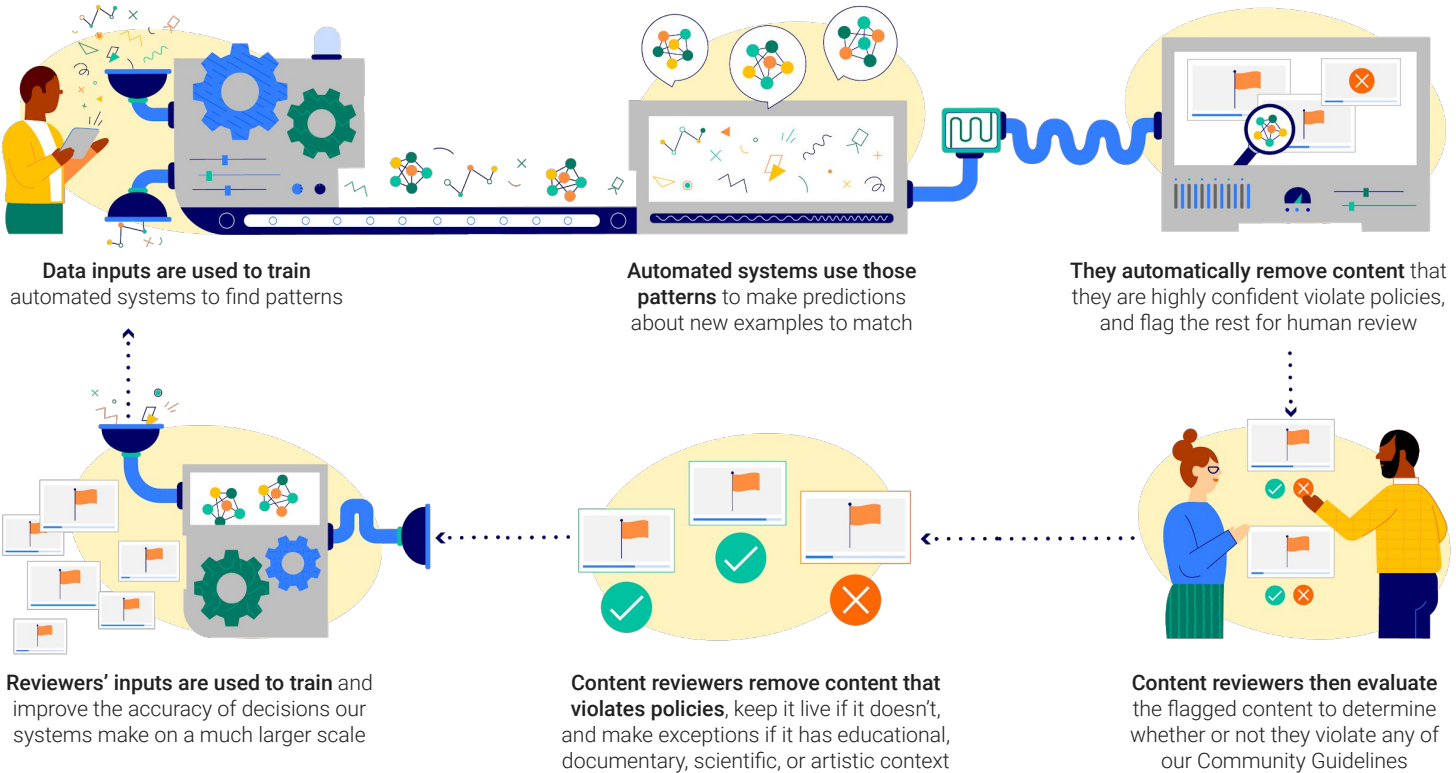
Who are content reviewers?

We have thousands of reviewers in time zones across the world who are fluent in multiple languages. They understand nuances of policy enforcement to evaluate flagged content.

Apply human judgement for complex situations where context is important, such as protecting content that has educational, documentary, scientific, or artistic context

Teach machine learning systems what we want to find by showing them examples of what we're looking for

How do content reviewers train machine learning systems?



9.3M

Total videos removed*

95%

Of removed videos were first detected by machine learning*

72%

Of videos removed for violating policies have less than 10 views*