

# Spam Filtering Overview

Last Modified on 02/01/2024 2:14 pm EST

## Table of Contents

- [How Spam Filtering Works](#)
- [Client Side Spam Filtering](#)
- [More Information](#)

About 50% of all email messages sent globally are spam. Email providers – including both Microsoft and Google – spend a lot of time, effort, and money making sure that most of that spam doesn't end up in your inbox (and that legitimate emails don't end up in your spam folder).

## How Spam Filtering Works

Broadly speaking there are two methods used to automatically flag spam messages: heuristics and neural networks. Each company has their own ways of using these methods, which they keep secret in an effort to thwart spammers.

### Heuristics

Heuristics is just a fancy way of identifying common characteristics of a thing and using those characteristics to categorize them. In this case, a lot of spam messages share qualities like:

- Country of origin
- Mismatched sender/reply to addresses
- Subject matter
- And more

Email systems keep track of these characteristics, or subscribe to a clearinghouse that provides lists of the characteristics to look out for, and flag messages with certain combinations as spam.

### Neural Networks

Google attributes its success in identifying and filtering spam to neural networks. Neural networks are, simply, a way in which one can train a computer, or a network of multiple computers, to make highly accurate decisions about what a thing might be.

This means that Google has a bunch of computers that have been trained, by looking at countless examples, to identify spam messages. But wait, you say, that sounds an awful lot like heuristics! Sort of, but not quite.

Think about when you send someone to the market with a shopping list. Heuristics is the kind of shopper who will stick exactly to the list. If the store doesn't have the exact yogurt you want, you aren't getting yogurt. If a spam message doesn't match something on the characteristics list, it isn't marked as spam. Neural Networks, however, are the sort of shopper who will see what you asked for but is willing to get something that meets the need even if it isn't an exact match. Neural Networks can flag things that perhaps aren't obviously spam according to some

checklist but meet the general spam profile that it has learned.

## Client Side Spam Filtering

[^Top](#)

Both of the methods described above are applied to your email messages by your email provider before the message ever gets delivered to you. They all happen on those remote email servers and as such are known as "server side" Spam filtering.

Client side filtering happens on your computer and in your email client (Outlook, for example). Most programs allow you to manually flag messages as spam, remove messages from your Junk folder, and create rules that will filter messages into Junk. These are all known as "client side" filters, and can be tweaked individually.

## More Information

[^Top](#)

- [Penn Email Routing](#)
  - [Why you might get Spam from Wharton addresses \(Email Spoofing\)](#)
  - [Creating an allowlist on Gmail \(Students\)](#)
  - [Creating Safe Senders in Microsoft O365](#)
-