SpamSieve 2.0 Manual

Michael Tsai
c-command.com

September 10, 2003

Contents

1 Introduction

1.1 What Is SpamSieve?

SpamSieve is an application that filters out unsolicited mass mailings, commonly known as "spam." Previously, most people just ignored spam messages or created simple rules in their e-mail clients to filter them out. In recent years and months, the spam problem has gotten worse. Today's spam is harder to detect, and there is more of it.

SpamSieve gives you back your inbox by bringing powerful Bayesian spam filtering to popular e-mail clients. It learns what your spam looks like, so it can block nearly all of it. It looks at your address book and learns what your good messages look like, so it won't confuse them with spam. Other spam filters get worse over time as spammers adapt to their rules; SpamSieve actually gets better over time as you train it with more messages. SpamSieve doesn't delete any messages—it only marks them in your e-mail client—so you'll never lose any mail. SpamSieve works with any number of mail accounts, of whatever types are supported by your e-mail software (e.g. POP, IMAP, Hotmail, AOL).

1.2 Identifying Spam

SpamSieve uses a statistical technique known as *Bayesian analysis*. For a more in-depth treatment of this technique applied spam, see this article by Paul Graham[1] and the papers it references. Bayesian spam filtering is highly accurate and adapts to new types of spam messages "in the field."

First, you *train* SpamSieve with examples of your good mail and your spam. When you receive a new message, SpamSieve looks at how often its words occur in spam messages vs. good messages. Lots of spammy words mean that the message is probably spam. However, the presence of words that are common in your normal e-mail but rare in spam messages can tip the scale the other way. This "fuzzy" approach allows SpamSieve to catch nearly every spam message yet produce very few false positives. (A *false positive* is a good message mistakenly identified as spam. Most users consider false positives to be much worse than *false negatives*—spam messages that the user has to see.)

Because you train SpamSieve with your own mail, you have full control. If SpamSieve makes a mistake, you can train it with the message in question so that in the future it will do better. Further, since spammers don't have access to the messages you trained SpamSieve with, they have no way of knowing how to change their messages to get through. Whereas other spam filters become less effective as spammers figure out their rules, *SpamSieve becomes more effective over time* because it has a larger corpus of your messages to work from.

1.3 Integration

Separate from the issue of identifying spam messages is the issue of how to prevent you from having to deal with them. There are basically six kinds of anti-spam software for doing this:

[1] http://www.paulgraham.com/spam.html

Challenge-Response Systems This software requires people who sends you mail to prove that they are human, and not an automated spam-sending program. After sending you a message, they get a reply asking them to complete a task that is easy for humans but hard for computers. Only then is the message passed on to you. This system is a nuisance for senders, delays your reception of the mail, and becomes impractical when sending messages to a group of people. Also, some legitimate messages *are* sent by programs, but the challenge-response system will treat them as spam.

Server-Side Filters This software runs on mail servers often filters out spam before you ever see it. This means that you do not have to download the spam messages that it catches. However, some spam messages may still get through, and, unless the filter is perfect, a few legitimate messages will not. These could be important messages, and you will never know that you lost them.

Server-Side Taggers This variant of server-side filters does not delete messages before you download them. Instead, you download every message and configure your e-mail client to move messages that were tagged by the filter into a separate spam folder. This eliminates the major disadvantage of server-side filters—lost messages—however this type of filter is generally not as accurate as the ones below, because it does not adapt to your own mail.

Client-Side Filters This software connects to your mail server to delete spam messages before your e-mail client can download them. This is a clunky approach: to catch all the spam messages, you have to run the program right before your regular e-mail program checks for mail. This is difficult to time properly if you check your mail often, and even so you may download some messages that weren't filtered. You will also download every good message twice. The anti-spam software may let you see the messages that it filtered out, so that you can verify that there were no false positives. However, you have to do this using its interface, not your e-mail program's (which is typically nicer). And if there was a false positive you then have to transfer it into your e-mail program so that you can file and reply to it.

Client-Side Proxies This is like a client-side filter except that the proxy downloads messages once and stores them locally. The e-mail client then "downloads" the good messages from the proxy. This addresses the timing and double-download problems of client-side filters, but interaction with the filter is still awkward because it happens outside your e-mail client. In addition, you lose some control over connections to the mail server and which messages are left on the server.

Client-Side Integrated This category includes SpamSieve and Apple Mail's built-in spam filter. Suspected spam messages are moved to a separate folder, which you can quickly scan at your leisure to make sure there are no false positives. The e-mail client downloads messages directly from the mail server, thus avoiding the problems of client-side filters and proxies. You can train the anti-spam software to improve its accuracy from inside your e-mail client, and accuracy is higher than with server-side filters because

the anti-spam software can learn from the messages that *you* receive. You can also control how the spam filter interacts with your regular mail sorting rules.

1.4 Main Features

- Powerful Bayesian spam filtering results in high accuracy and almost no false positives. It adapts to the mail that *you* receive to get even better with time.
- Integrates with your e-mail client for a superior user experience.
- Integrates with the Mac OS X Address Book so that messages from friends and colleagues are never marked as spam.
- Automatically maintains a blocklist so that it can instantly adapt to spam messages sent from particular addresses, and catch 100% of them.
- Automatically maintains a whitelist to guarantee that messages from particular senders are never marked as spam, without cluttering your address book with these addresses.
- Can honor Habeas[2] headers warranting that a message is not spam, as well as the "ADV" subject tag indicating that a message *is* spam.
- Many spammers encode the contents of their messages so that filters cannot see the incriminating words they contain. SpamSieve can decode and look inside these messages. Optionally it can mark them all as spam, on the theory that legitimate senders do not try to obscure their messages.
- SpamSieve keeps track of how accurate it is, how many good and spam messages you receive, and how these numbers change over time.
- Turn off new-mail notification in your e-mail client, and let SpamSieve notify you only when you receive non-spam messages.
- The corpus window and log let you see how each spam message was caught.

2 Installing and Updating

2.1 Requirements

SpamSieve has been developed and tested on Mac OS X 10.1.5 and 10.2.6. It is designed to work with the following e-mail clients:

- Emailer 2.0v3[3], previously available from Claris
- Entourage 9.0.1[4] and later from Microsoft

[2] http://www.habeas.com/partner.php?affil=spamsieve
[3] http://www.fogcity.com
[4] http://www.microsoft.com/mac/entouragex/default.asp?navindex= s4

- Eudora 5.2[5] and later (6.0 recommended) from Qualcomm
- Mailsmith 1.5[6] and later (2.0.1 recommended) from Bare Bones Software
- PowerMail 4.0[7] and later from CTM Development

2.2 Updating From a Previous Version

SpamSieve 2.0 will automatically read the corpus and statistics from previous versions. You can simply replace the old application file with the new one.

SpamSieve now works with Eudora 6's spam filtering menu commands. Eudora users may want to replace the SpamSieve Eudora Helper with the new SpamSieve Eudora Plug-In. See Section 4.3.1.

This version of SpamSieve contains updated AppleScripts for improved integration with e-mail client software. If you want to take advantage of these improvements (listed in Section 10), you will need to replace your existing scripts with the new ones. You may need to reconfigure the mail rules/filters in your e-mail client so that they reference the new scripts. If you have modified the existing scripts, be sure to transfer the modifications to the new scripts.

2.3 Installing SpamSieve

Double-click the `SpamSieve-2.0.dmg` file to mount or expand the SpamSieve disk image. Then move the SpamSieve application to your `Applications` folder. There's no need to copy this manual to your hard disk. A copy of it is built into SpamSieve, and you can access it by choosing **SpamSieve Manual (PDF)** from the **Help** menu.

Next, follow the instructions in Section 4 for setting up SpamSieve to use it with your e-mail client.

After setting up SpamSieve, you will need to train it with examples of your spam messages and good messages (Section 3).

2.4 Uninstalling SpamSieve

To uninstall SpamSieve, delete any rules that you created for it in your e-mail client. You can also delete the AppleScripts and/or plug-in that you installed. If you are using the SpamSieve Eudora Helper, delete it, and also run the Uninstall Eudora Helper program that came with SpamSieve. You can also delete the SpamSieve application and its data files, which are stored in `~/Library/Application Support/SpamSieve`.

[5] http://www.eudora.com/mac
[6] http://www.barebones.com/products/mailsmith.html
[7] http://www.ctmdev.com/powermail4.shtml

3 Training SpamSieve to Recognize Your Spam

Before you can use SpamSieve, you must give it some examples of messages you consider to be spam, and ones which you do not. You do this by selecting some messages and then telling SpamSieve to add them to its corpus. For the details of how to do this, see Section 4. For now, what's important is that you will train SpamSieve with both good messages and spam messages.

In general, the more messages you train SpamSieve with, the better its accuracy will be. For best results, you should train it with *at least* 600 messages. Training SpamSieve with more of one type of message will bias its predictions to that type. For instance, a corpus with mostly good messages will make SpamSieve conservative about identifying spam, leading to more false negatives. Most users find that filling the corpus with two to four times as many spam messages as good ones (as shown at the bottom of the **Statistics** window) produces a low level of false negatives, while keeping false positives rare or non-existent. For this reason, you should *not* add all your good messages to the corpus when you first install SpamSieve, because you likely do not have enough saved spam messages to compensate.

If SpamSieve marks a good message as spam, you should add the message to SpamSieve as a good message. This lets SpamSieve know that it made a mistake, and also adds the message to the corpus to improve future accuracy. Likewise, if SpamSieve marks a spam message as good, you should add the message to SpamSieve as a spam message. *If you do not correct SpamSieve when it makes mistakes, its accuracy will deteriorate over time.*

If you make a mistake and tell SpamSieve that a message is spam when it is actually good (or vice-versa), simply correct yourself as you would correct SpamSieve. That is, if the message is good, add the message as good; if it is spam, add it as spam. SpamSieve will "undo" the previous, incorrect, addition to the corpus.

To improve SpamSieve's accuracy, it is important to train it with new messages as you receive them. When they first install SpamSieve, most users have many good messages to train it with, but few spams. For this reason, SpamSieve is set to automatically train itself with spam messages as you them. It does not automatically train itself with good messages, on the assumption that you have already trained it with plenty of good mail. Both of these settings may be modified in the preferences. As a result, after training SpamSieve for the first time, you only need to train it to correct mistakes.

4 Using SpamSieve With Your E-Mail Client

4.1 Emailer

4.1.1 Installing

Copy the files from SpamSieve's `For Emailer Users` folder into Emailer's `AppleScripts` folder:

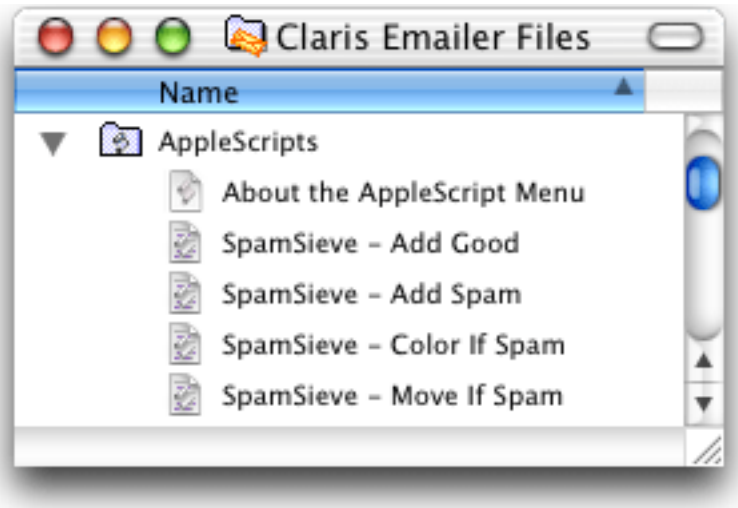

You may need to quit and re-launch Emailer in order for it to notice that you have installed the SpamSieve AppleScripts.

If you want SpamSieve to color messages that it thinks are spam, set up a mail action in Emailer that looks like this:

If, instead, you want SpamSieve to move suspected spam messages to a **Spam** folder (that it creates automatically), set up a mail action in Emailer that looks like this:

SpamSieve will now automatically color or move new spam messages that you receive, depending on which mail action you set up.

4.1.2 Training SpamSieve

To train SpamSieve with spam messages, select one or more of them and then choose **SpamSieve - Add Spam** from Emailer's **Scripts** menu. To train SpamSieve with good messages, select one or more of them and then choose **SpamSieve - Add Good** from Emailer's **Scripts** menu.

4.1.3 Manually Processing Messages

To manually ask SpamSieve to color or move messages that it thinks are spam, select the messages and choose **SpamSieve - Color If Spam** or **SpamSieve - Move If Spam** from Emailer's **Scripts** menu.

4.2 Entourage

4.2.1 Installing

Go to Entourage's **Tools** menu select **Junk Mail Filter**. Make sure the Junk Mail Filter is disabled. Then open the Mailing List Manager, also in the **Tools** menu, and make sure that there are no items listed there. Both of these tools can interfere with SpamSieve.

Locate the `Entourage Script Menu Items` folder inside the `Microsoft User Data` folder (which is probably in your `Documents` folder). Copy the files from SpamSieve's `For Entourage Users` folder into the `Entourage Script Menu Items` folder:

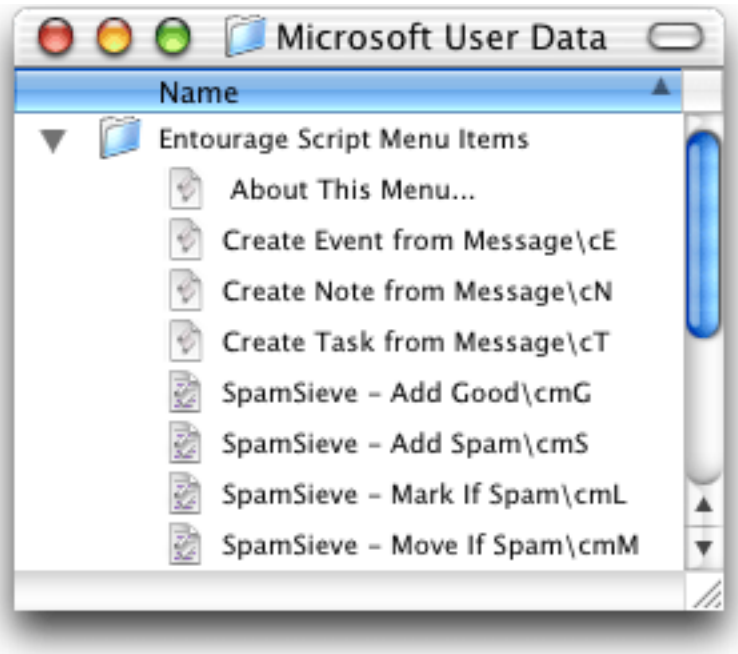

You may need to quit and re-launch Entourage in order for it to notice that you have installed the SpamSieve AppleScripts. Then you should see four SpamSieve items in Entourage's **Scripts** menu:

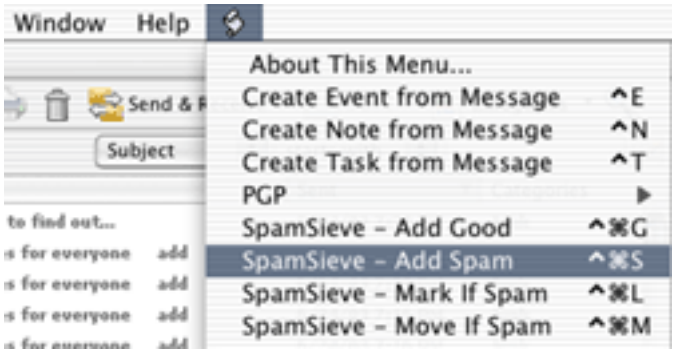

If you do not see the SpamSieve scripts in the menu, it may be because you have more than one `Entourage Script Menu Items` folder. Try holding down the Option key as you choose one of the items from Entourage's **Scripts** menu. This will open a script file in Script Editor. Then hold down the Command key as you click on the name of the script in the window's title bar. This will show you the location of the folder where Entourage is looking for its scripts:

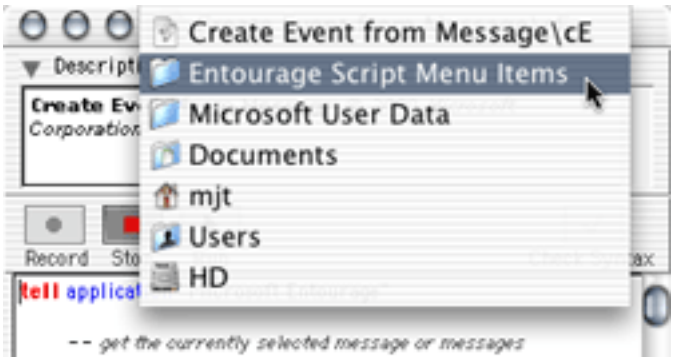

If you want SpamSieve to color messages that it thinks are spam, set up a mail rule in Entourage that looks like this:

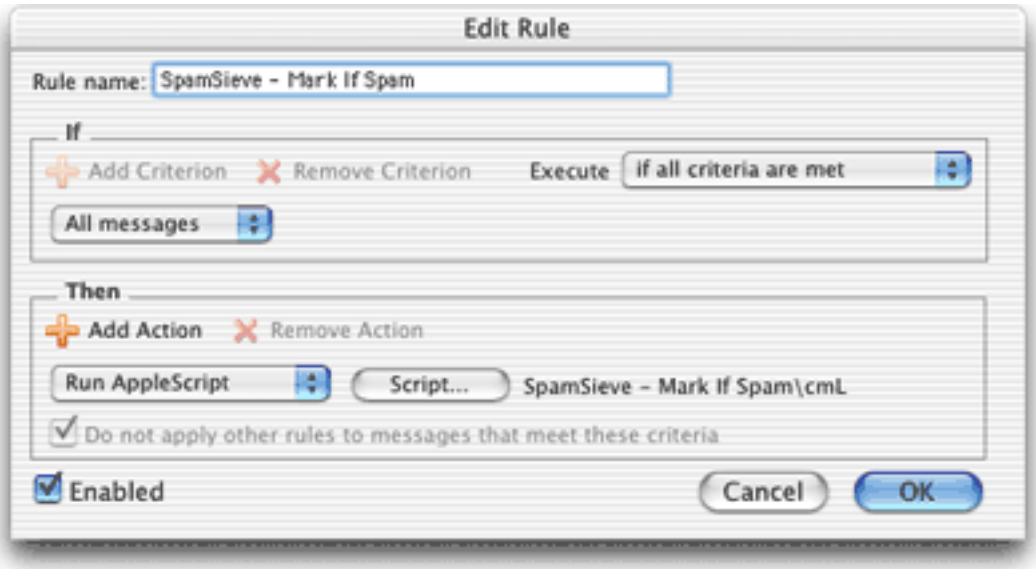

To do this, choose **Rules** from Entourage's **Tools** menu. Click on the tab corresponding to the type of account you have (e.g. POP). If you have more than one kind of account, you will need to create a rule for each account type. Click the **New** button. Type a name for your rule. Then click just to the left of **Change status** to select the first action. Click **Remove Action**. Click on the menu that says **Set category** and select **Run AppleScript**. Then click the **Script**... button and select the `SpamSieve - Mark If Spam` file. The rule window should now look like the above screenshot.

This will make SpamSieve mark new spam messages that you receive by changing their category and color. If, instead, you want SpamSieve to move suspected spam messages to a **Spam** folder (that it creates automatically), set up the mail rule to use the **SpamSieve - Move If Spam** script instead:

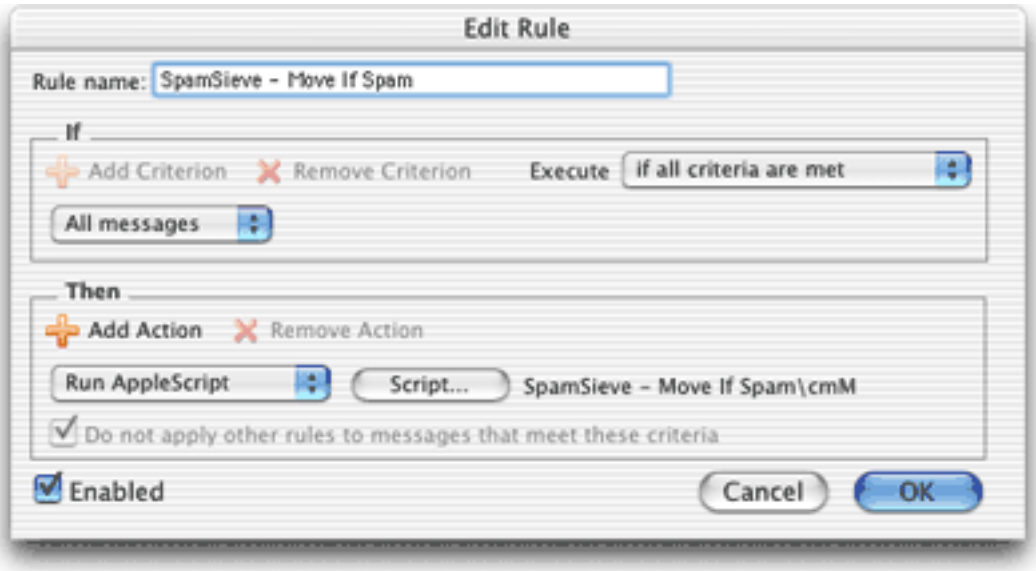

As a third option, you can instead create the rule such that SpamSieve never marks messages from people in your Entourage address book as spam. Do to this, set up the **If** section of the rule like so:

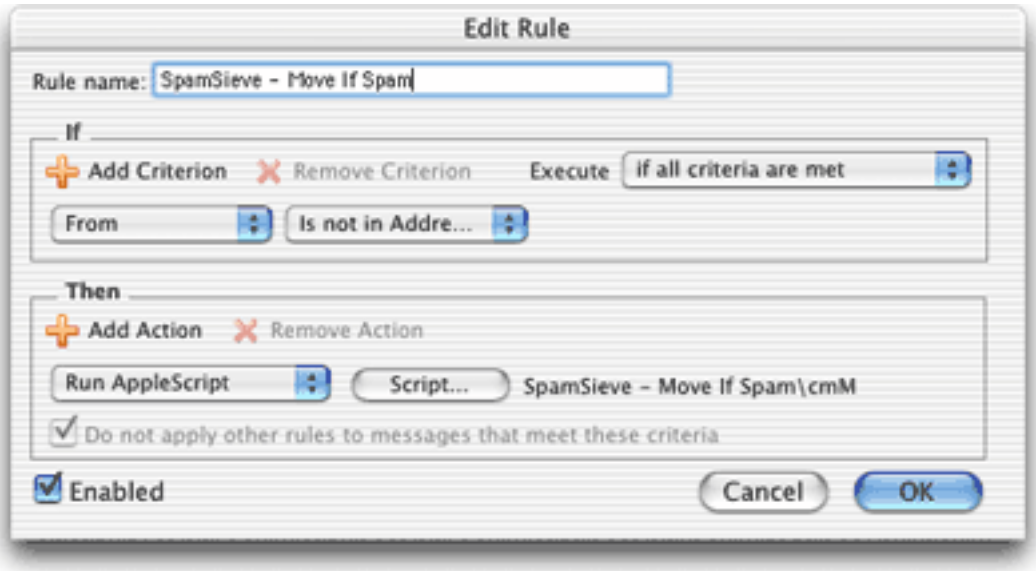

It is important that you create the rule exactly as shown. Do not add additional actions below the action that runs the AppleScript. Such actions would apply to all messages, which is probably not what you want. To customize what Entourage does when SpamSieve finds a spam message, you need to edit the AppleScript rather than the rule.

4.2.2 Training SpamSieve

To train SpamSieve with spam messages, select one or more of them and then choose **SpamSieve - Add Spam** from Entourage's **Scripts** menu. To train SpamSieve with good

messages, select one or more of them and then choose **SpamSieve - Add Good** from Entourage's **Scripts** menu.

Note that due to the way Entourage's AppleScript interface works, you may not be able to train SpamSieve by selecting messages in custom mail views. Instead, select the messages in their actual folders.

Make sure that you correct SpamSieve's mistakes by using the commands in the **Scripts** menu—do not click the underlined blue text to indicate that a message is not spam.

4.2.3 Manually Processing Messages

To manually ask SpamSieve to mark or move messages that it thinks are spam, select the messages and choose **SpamSieve - Mark If Spam** or **SpamSieve - Move If Spam** from Entourage's **Scripts** menu.

4.2.4 IMAP Accounts

Entourage does not support moving IMAP messages via AppleScript, so if you use IMAP the **SpamSieve - Move If Spam** script will not move spam messages into your **Spam** folder.

The `SpamSieve - Add Spam` and `SpamSieve - Move If Spam` scripts contain an experimental workaround for moving IMAP messages. You can enable this by editing the scripts in Script Editor and changing the `false` after `tryToMoveIMAPMessages` to `true`. However, some users have found that this exposes a bug in Entourage, causing it to crash.

If you do not require IMAP, you can try creating a POP account in Entourage and re-entering your account information. Many IMAP accounts also work via POP, and this will allow SpamSieve to move the messages that it thinks are spam.

Alternatively, you can create an Entourage rule that moves messages that SpamSieve has marked as junk into another folder. After receiving mail, manually apply this rule to the messages in your IMAP account.

Yet another alternative is to use a custom view. Make a custom view of your IMAP account that looks for messages that are not junk. In this way, you can view your good messages without being distracted by spams.

4.2.5 Advanced Rules

Two complications are that once an Entourage rule runs an AppleScript or moves a message, it cannot apply any more rules to that message. Here are some ways to work around this.

One option is to order your rules so that Entourage applies the SpamSieve rule after all your other rules. You can change the order of the rules by choosing **Rules** from Entourage's **Tools** menu and dragging the rules in the list to change their order. With this approach, you can filter your good messages into folders however you want. Any mail that is not moved into another folder will remain in your inbox. Then, the SpamSieve rule will either mark the spams as junk or move them to a **Spam** folder. The disadvantage to this approach is that SpamSieve cannot catch any spams among the messages that were moved by your other rules.

Another option is to add the **Run AppleScript** action to each rule that moves messages. For instance, suppose you have a rule that moves all the messages from your **Work Account** account into a **Work** folder. You could set up the rule like this:

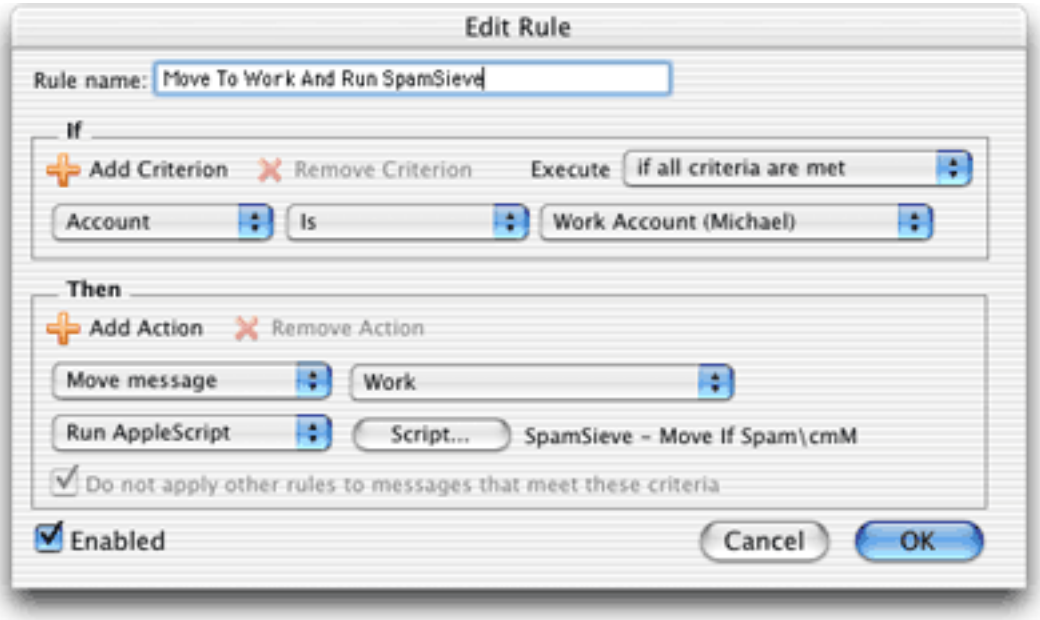

Now, messages sent to that account will be moved to the **Work** folder. Spam messages sent to that account will be moved to the **Spam** folder.

You can add the SpamSieve AppleScript action to every rule that moves messages and also to a "catch-all" rule that applies to messages that aren't moved. Then SpamSieve will be able to filter all the messages that you receive.

Please contact spamsieve@c-command.com[8] if you have trouble setting up Entourage to filter messages the way you want.

4.3 Eudora

4.3.1 Installing

If you are using Eudora 5.2, please see the instructions in Appendix B. If you previously used SpamSieve 1.x with Eudora, run the `Uninstall Eudora Helper` applet in the `For Eudora Users` folder.

If you are using Eudora 6, locate the Eudora application file and choose **Get Info** from the Finder's **File** menu. Expand the **Plug-ins** pane and click **Add**... Then locate the `SpamSieve Eudora Plug-In` file in SpamSieve's `For Eudora Users` folder. Finally, uncheck the **SpamHeaders** and **SpamWatch** plug-ins so that they do not interfere with SpamSieve. When you start up Eudora, you should see SpamSieve listed in the **About Message Plug-ins**... window that is accessible from the **Eudora** menu.

[8] mailto:spamsieve@c-command.com

Now Eudora will use SpamSieve to filter all incoming messages. It will move the spam messages to the **Junk** mailbox.

Normally, Eudora will launch the SpamSieve application when new messages arrive or when you train SpamSieve from inside Eudora. However, on some machines, it will not launch SpamSieve automatically. In this case, you must manually open the SpamSieve application when you want Eudora to filter spam messages.

4.3.2 Training SpamSieve

To train SpamSieve with spam messages, select one or more of them and then choose **Junk** from Eudora's **Message** menu. To train SpamSieve with good messages, select one or more of them and then choose **Not Junk** from Eudora's **Message** menu.

4.3.3 Setting Options

Eudora applies SpamSieve to all incoming messages. The **Junk Mail** area of Eudora's preferences lets you customize how Eudora interacts with SpamSieve. Note that the **Junk Threshhold** [**sic**] slider will have no effect because SpamSieve always gives Eudora scores that are exactly 0 or exactly 100. Instead of using this slider, you should use the one in the **Bayesian** tab of SpamSieve's preferences.

The **Junk Extras** area of Eudora's preferences lets you control some additional settings, such as whether junk messages are removed from the mail server. If you do not see the **Junk Extras** icon in Eudora's preferences, you can make it available by enabling Eudora's Esoteric Settings plug-in. To do this, locate the Eudora application file and choose **Get Info** from the Finder's **File** menu. Expand the **Plug-ins** pane and make sure there is a check next to **Esoteric Settings 6.0**.

4.4 Mailsmith

4.4.1 Installing

Mailsmith 2.0 and later feature direct integration with SpamSieve. This is more convenient and easier to use than the script- and filter-based integration that was necessary when using previous versions of Mailsmith (see Appendix A). You can enable SpamSieve simply by clicking the checkbox in the **Spam Handling** pane of Mailsmith's preferences. For more information about using SpamSieve with Mailsmith, please see Chapter 8 of the Mailsmith User Manual.

4.4.2 Training SpamSieve

To train SpamSieve with spam messages, select one or more of them and then choose **Mark as Spam** from Mailsmith's **Message** menu. To train SpamSieve with good messages, select one or more of them and then choose **Mark as Non-Spam** from Mailsmith's **Message** menu.

4.4.3 Setting Options

You can configure how Mailsmith and SpamSieve work together from the **Spam Handling** pane of Mailsmith's preferences. Checking the **Train the Spam Detector** checkbox here is equivalent to checking *both* auto-training checkboxes in SpamSieve's preferences. It is recommended that you uncheck **Train the Spam Detector** once you have trained SpamSieve with a few hundred messages of each type.

4.4.4 Identifying Spam Messages

When using SpamSieve with Mailsmith, Mailsmith tags messages using the **Is Spam** and **Is Not Spam** properties. Although you can use Mailsmith's **Advanced Query** feature to search on these properties, they are otherwise not visible in the user interface. Therefore, you should mark spam messages in a visible way, either by letting Mailsmith move them to a separate mailbox, or by setting up a filter to change the messages' labels based on their **Is Spam** and **Is Not Spam** properties. Otherwise, you will not be able to correct SpamSieve's mistakes to improve its accuracy.

4.5 PowerMail

4.5.1 Installing

Copy the files from SpamSieve's `For PowerMail Users` folder to the `PowerMail Scripts` folder inside the `PowerMail Files` folder. The `PowerMail Files` folder is probably located in your `Documents` folder.

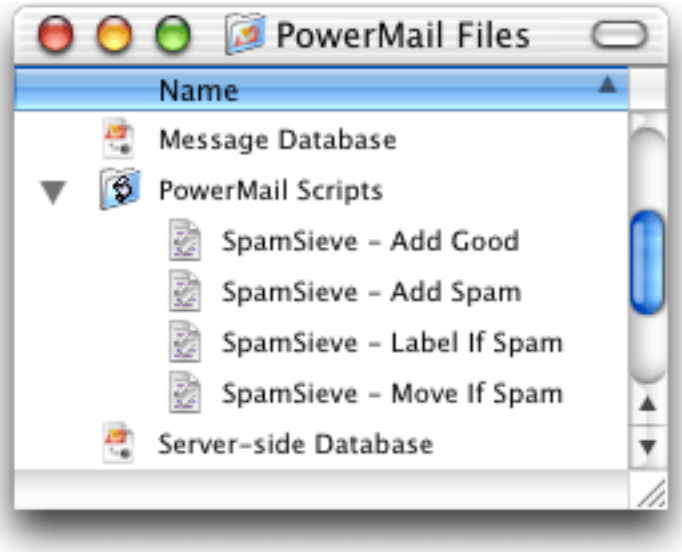

If you want SpamSieve to label messages that it thinks are spam, set up a filter in PowerMail that looks like this:

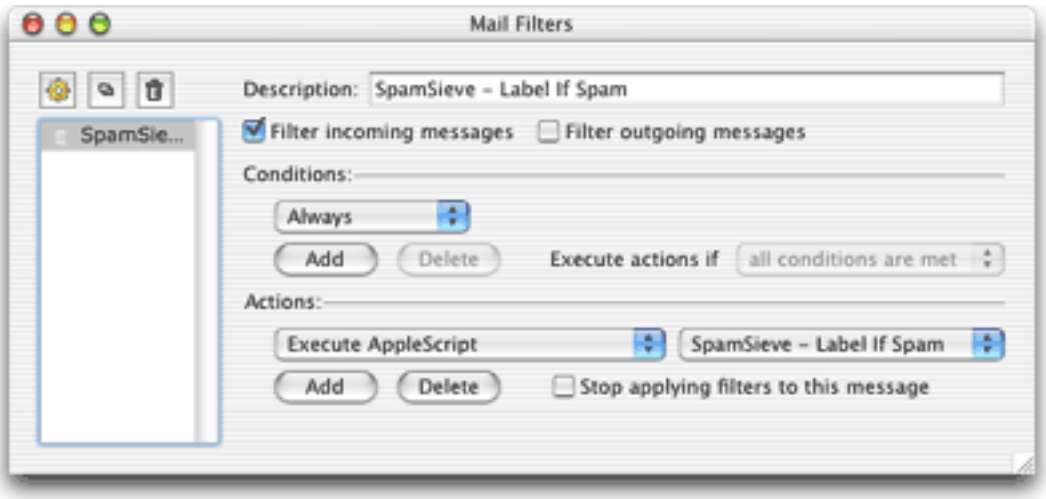

If, instead, you want SpamSieve to move suspected spam messages to a **Spam** folder (that it creates automatically), set up the filter to use the `SpamSieve - Move If Spam` script instead.

SpamSieve will now mark or move new spam messages that you receive.

4.5.2 Training SpamSieve

To train SpamSieve with spam messages, select one or more of them and then choose **SpamSieve - Add Spam** from PowerMail's **Scripts** menu. To train SpamSieve with good messages, select one or more of them and then choose **SpamSieve - Add Good** from PowerMail's **Scripts** menu.

4.5.3 Manually Processing Messages

To manually ask SpamSieve to label or move messages that it thinks are spam, select the messages and choose **SpamSieve - Label If Spam** or **SpamSieve - Move If Spam** from PowerMail's **Scripts** menu.

4.5.4 IMAP Accounts

PowerMail does not support moving IMAP messages via AppleScript, so if you use IMAP the **SpamSieve - Move If Spam** script will not move spam messages into your **Spam** folder. IMAP users should create the rule that uses the labelling script.

4.6 General Filtering Advice

Some of the e-mail clients that SpamSieve supports let you control the order in which the rules (a.k.a. filters or mail actions) that you have created process mail. How you order the SpamSieve rule is up to you. If you get a lot of spam that matches the rules you use to organize your mail, you might want to run the SpamSieve rule first. This will allow it to find spam among all your messages. If you would rather deal with spam manually than have

any false positives, then you might want to run the SpamSieve rule last, after all your other rules have been given a chance to match and file away messages from known senders. Be sure to check the SpamSieve preferences for additional filtering options.

4.7 Customizing the AppleScripts

4.7.1 Moving Spam Messages After Training

Normally when you add spam messages to the corpus, SpamSieve moves the messages to a **Spam** folder in your e-mail client. If you prefer that it not move the messages, you can do the following. Use Script Editor to open the `SpamSieve - Add Spam` AppleScript for your e-mail client. (Mailsmith and Eudora users do not have such a script.) The first line of the script contains the text:

```
property shouldMoveSpam : true
```

Change the `true` to `false` so that you have:

```
property shouldMoveSpam : false
```

Then save the script and close its window.

There are other aspects of the SpamSieve AppleScripts that you can easily customize by editing the scripts in Script Editor. Take a look at the different properties listed at the top of each script.

4.7.2 Compatibility Notes

Some Mac e-mail clients only work with the older resource-fork format for compiled AppleScripts. This format is no longer the Script Editor default in Mac OS X 10.2 and later. To save scripts in the older format, make sure that you edit the supplied scripts (or copies of them) rather than creating new compiled script files from inside Script Editor. If you create a new file, Script Editor will save it in the, incompatible format.

The Script Editor 2.0 beta cannot edit the `SpamSieve Eudora Helper` applet. To customize the applet, use Script Editor 1.9.

4.7.3 Integrating With Other Applications

SpamSieve's interface for integrating with third-party mail and news programs is completely open. It is possible to add support for additional programs simply by writing some AppleScripts. SpamSieve's AppleScript dictionary contains some basic information about the supported commands. However, there are some subtle, but important, points that are not discussed in the dictionary's documentation. If you would like to connect an application to SpamSieve, please contact spamsieve@c-command.com[9] so that I may assist you.

[9] mailto:spamsieve@c-command.com

5 Menus

5.1 The File Menu

5.1.1 Import Corpus...

This imports the words in a corpus that was exported in XML format. This is the same format used by SpamSieve 1.x. Importing a corpus merges it with the active corpus. To replace the active corpus with the one you are importing, use the **Reset Corpus**... command before importing.

5.1.2 Export Corpus...

This exports the active corpus to XML format. You might do this in order to use the corpus on another machine.

5.1.3 Import Messages...

This imports messages stored in mbox format. You can select whether the mbox file contains good messages or spam messages. To replace the active corpus with the messages you are importing, use the **Reset Corpus**... command before importing.

5.1.4 Import Seed Spam...

Some users do not have many saved spam messages with which to train SpamSieve. This command adds about 1400 spam messages (from a public archive) to the corpus in order to jump-start spam recognition. You still need to train SpamSieve with your own good messages, however. Note that accuracy will be better if you train SpamSieve with your own spam rather than the seed spam. Thus, you should only import seed spam if you have few saved spam messages and do not receive many new ones per day.

5.2 The Filter Menu

5.2.1 Show Corpus

This command opens the **Corpus** window so that you can examine the words that SpamSieve has found in your e-mails. You can click on the name of a column to sort by that column. Click again on the column to reverse the sort direction. The meanings of the columns are as follows:

Word A word in the corpus.

Spam The number of times the word has occurred in spam messages.

Good The number of times the word has occurred in good messages.

Total The total number of times the word has occurred.

Prob. The probability that a message is spam, given that it contains the word (and in the absence of other evidence).

Last Used The date that the word was added to the corpus, or the date that it last appeared in a received message (whichever is later).

You can copy the selected rows to the clipboard or drag and drop them into another application.

With the window sorted by **Word**, you can type the first few letters of a word to locate that word in the corpus. Similarly, you can sort by one of the other columns and type a number to locate the first word whose value for the sorted column matches the number you typed.

You can edit the spam and good counts associated with a word by double-clicking on the number in the **Spam** or **Good** column. Changing the numbers for important words can greatly affect SpamSieve's accuracy, so you shouldn't make changes without good reason.

You can remove words that you don't want in the corpus by selecting them and pressing Delete.

5.2.2 Prune Corpus...

This command removes from the corpus words that have not been used in a set number of days. This can decrease SpamSieve's memory use and increase its speed. However, pruning can reduce SpamSieve's accuracy, so you should only prune if you find that SpamSieve is using too many system resources.

5.2.3 Reset Corpus...

This command removes all the words and messages from the corpus. This will enable you to retrain SpamSieve from scratch, and SpamSieve will let you use your old messages in the retraining.

5.2.4 Show Statistics

This command opens the **Statistics** window, which displays the following information:

Good Messages The number of non-spam messages that SpamSieve has filtered.

Spam Messages The number of spam messages that SpamSieve has filtered.

False Positives The number of good messages that SpamSieve identified as spam.

False Negatives The number of spam messages that SpamSieve identified as good.

% Correct The percent of messages that SpamSieve identified correctly.

Good Messages The number of non-spam messages that are used to identify spam messages.

Spam Messages The number of spam messages that are used to identify spam messages.

Total Words The total number of unique words in the corpus.

You can copy all the statistics to the clipboard using the **Copy** command in the **Edit** menu or by holding down the Control key and clicking in the window.

Normally, SpamSieve shows statistics for all the messages that it processed since it was installed. Because the accuracy and the number of messages you receive change with time, you may wish to only see recent statistics. Click the **Set**... button at the bottom of the window to control which old statistics are hidden from view. You can edit the date and time shown in the sheet, or enter an entirely new date. SpamSieve will accept dates specified in natural language, such as "last Sunday at dinner."

Note that all messages processed by SpamSieve 1.2.2 (and earlier) are recorded on the date that you first launched SpamSieve 1.3. There is no fine-grain control over which of these messages are displayed in the statistics.

5.2.5 Open Log

SpamSieve keeps a log of messages that it has filtered, words that it has used to decide whether the messages were spam, messages you have added to the corpus, and any errors that have occurred. This command opens the log file so that you can look at it. Normally, there is no reason (aside from curiosity) to look at the log file. However, if you believe SpamSieve is not working as it should, the log file provides useful information about what SpamSieve has done. If you find that the log file is taking up too much disk space, you can delete it at any time. SpamSieve will then start a new log file as needed.

5.2.6 Show Blocklist

This opens the **Blocklist** window. Messages sent from addresses on the blocklist will always be marked as spam. The meanings of the columns are as follows:

Date The date that the rule was added to the blocklist.

Sender The sender address to block messages from.

√ If this is checked, the rule is enabled. Disabled rules do not block any messages.

Hits The number of spam messages that the rule blocked.

Training SpamSieve with a spam message adds its sender to the blocklist. You can delete a rule from the blocklist by selecting it and pressing Delete. You can copy the selected rows to the clipboard or drag and drop them into another application. You can type the first few letters of an address to locate that sender in the list.

5.2.7 Show Whitelist

The whitelist works the same way as the blocklist except that messages sent from addresses on the whitelist are never considered to be spam.

6 Preferences

6.1 Filters

6.1.1 Order

SpamSieve uses a variety of filters to determine whether messages are spam or good. It consults the filters in the order listed in this window. When a filter decides that the message is good or spam, SpamSieve stops moving down the list. Thus, the order of the filters makes a difference. You can see from the order that if a message's sender is on the whitelist, it will be marked as good even if the Bayesian classifier would have predicted it to be spam. Normally this is what you want; the point of a whitelist is that you can be sure that certain messages will *never* be marked as spam.

6.1.2 Check for message in corpus

SpamSieve learns as you train it, but training is not instant. Training SpamSieve with a message will not necessarily give it enough information to classify that message correctly based only on the words in the message. However, once you have added a message to the corpus, SpamSieve *knows* whether it is good or spam, even though it might not make the correct prediction based on word probabilities. This option causes SpamSieve to see if it knows whether a message is good or spam before trying to calculate its spam probability. If SpamSieve has seen the message before, it will always classify it correctly. You can disable this option if you want to see what SpamSieve would have predicted if it did not already know whether the message was good or spam.

6.1.3 Use Mac OS X Address Book

With this option enabled, SpamSieve will never predict a message to be spam if its sender's e-mail address is in the system address book. This feature requires Mac OS X 10.2 or later.

You can add addresses to the system address book using the Address Book application (located in the `/Applications` folder), or directly from an e-mail client that supports the system address book.

Mailsmith and PowerMail users should be sure to enable the option to use Apple's Address Book. Entourage users may prefer to use their address book as a whitelist instead of Apple's. This is described in Section 4.2.1.

6.1.4 Exclude my addresses

Enable this option so that spam messages with your own return address are not marked as good. Disable it if you send yourself messages and want to make sure that they are never marked as spam.

6.1.5 Use SpamSieve whitelist

Enable this option so that messages sent from addresses on the SpamSieve whitelist are never marked as spam.

6.1.6 Use SpamSieve blocklist

Enable this option so that messages sent from addresses on the SpamSieve blocklist are always marked as spam.

6.1.7 Honor Habeas headers

The Habeas[10] service lets legitimate users add special headers to their messages to indicate that they are not spam. This option makes SpamSieve trust those headers, so that it never marks a message as spam if the headers are present. Habeas takes legal action against spammers who use the headers.

6.1.8 "ADV" messages are spam

This option causes SpamSieve to always mark messages as spam if they contain some variant of "ADV" at the start of the subject line. The "ADV" marker is used by some commercial bulk mailers.

6.1.9 Encoded HTML mail is spam

Many spammers encode the contents of their messages with base-64 so that filters cannot see the incriminating words they contain. SpamSieve can decode and look inside these messages. This option causes it to mark *all* such as spam, regardless of their contents, on the theory that legitimate senders do not try to obscure their messages.

6.1.10 Use Bayesian classifier

This enables SpamSieve main spam detector, which uses the corpus and word probabilities to identify spam messages.

6.1.11 Save false negatives to disk

False negatives are spam messages that SpamSieve didn't catch. This option causes SpamSieve to save such messages for later analysis. Clicking the **Show** button opens the folder containing the saved messages. You can e-mail this folder, or selected files from it, to spamsieve-fn@c-command.com[11]. By looking at the messages that SpamSieve missed, I can improve its algorithms to catch such messages in the future. Note that enabling this option will slow down SpamSieve's processing.

[10] http://www.habeas.com/partner.php?affil=spamsieve
[11] mailto:spamsieve-fn@c-command.com

6.2 Notification

6.2.1 What Is Notification?

All e-mail clients can notify you when you receive new messages, but most will notify you even when all the new messages are spam. If your e-mail client is not savvy in this way, you can turn off its notification and let SpamSieve notify you only when there are new good messages.

6.2.2 Play sound

This makes SpamSieve play a sound when new good messages are received. To add a sound to the menu, copy the sound file to the `Sounds` folder in your `Library` folder.

6.2.3 Bounce Dock icon

This makes SpamSieve bounce its Dock icon once when new good messages are received.

6.2.4 Keep bouncing until clicked

You might not be looking at the Dock icon when it first bounces, so this makes SpamSieve continue bouncing its Dock icon until you click on it.

6.2.5 Show count in Dock icon

This option makes SpamSieve show the number of new good messages in its Dock icon. If there are no new good messages, SpamSieve will not show any number (rather than showing zero). Clicking the Dock icon or training SpamSieve with a message will reset the count.

6.3 Training

6.3.1 Allow duplicates in corpus

If you allow duplicate messages in the corpus, adding the same message twice will increase the counts for the words in that message. If you do not allow duplicate messages, the second and subsequent times you try to add a message will have no effect. By default, duplicate messages are not allowed in the corpus. This is nice because it means that you do not have to remember which messages you have already added; accidentally adding the same message more than once will not skew the data that you are providing to SpamSieve. If you wish to intentionally skew the data, you can check one or both boxes to allow duplicates.

6.3.2 Auto-train

These options cause SpamSieve to automatically train itself with newly received messages based on their predicted categories. Thus, after the initial training you would only need to train SpamSieve to correct its mistakes; it would automatically learn from all the other new messages.

Note that with either of these options enabled it is imperative that you correct SpamSieve when it makes a mistake; otherwise it will be making predictions based on incorrect information.

Auto-training with good messages tends to prevent false positives while slightly increasing false negatives. Auto-training with spam messages tends to prevent false negatives while slightly increasing false positives. Most users will train SpamSieve with plenty of good messages right after installing it, but they may not have many spam messages on hand for the initial training. Thus, it is recommended that you have SpamSieve only auto-train with spam messages to its corpus.

SpamSieve works best when trained only on misclassified messages. Once it has attained a good accuracy, you should turn off auto-training all together.

6.3.3 Train SpamSieve whitelist

With this option enabled, training SpamSieve with a good message will add the message's sender to SpamSieve's whitelist. Training SpamSieve with a spam message will disable the sender if it appears in the whitelist.

Example: You receive an Amazon order receipt and train SpamSieve with it as a good message. This puts `auto-confirm@amazon.com` on the whitelist so that future order receipts are always marked as good. A spammer might decide that `auto-confirm@amazon.com` would make a good fake return address. If you receive such a spam, SpamSieve would mark it as good because the sender was on the whitelist. If you then tell SpamSieve that the message is spam, it will disable the whitelist rule so that it can catch future spam messages with that return address.

6.3.4 Train SpamSieve blocklist

With this option enabled, training SpamSieve with a spam message will add the message's sender to SpamSieve's blocklist. Training SpamSieve with a good message will disable the sender if it appears in the blocklist.

6.3.5 Train Bayesian classifier

With this option enabled, training SpamSieve with a message will add the words from that message to SpamSieve's corpus. It is highly recommended that you train the Bayesian classifier, as this is how most spam messages are caught.

6.4 Bayesian

This slider lets you adjust SpamSieve's bias. The bias controls how aggressive SpamSieve is at catching spam. When SpamSieve is more aggressive, it is better at catching spam messages that share some characteristics with your good mail. When SpamSieve is more conservative, it will be better at marking borderline messages such as order confirmations and press releases as good. This setting is very powerful, and most users should stick to the

middle range. It is also not a substitute for training SpamSieve. Only change the bias if SpamSieve is consistently making errors in the same direction.

7 Frequently Asked Questions

7.1 How can I hide SpamSieve's Dock icon?

The easiest way is to use the free Dockless[12] utility. You'll need to make the Dock icon visible again in order to configure SpamSieve's preferences or view the statistics.

7.2 What information should I include when I report a problem?

If you are reporting a problem with SpamSieve's accuracy, open the **Statistics** window and choose **Copy** in the **Edit** menu so that you can paste all your statistics into the message. Also, use the **Open Log** command in the **Filter** menu and include any relevant entries from the log.

7.3 Where can I download the older Mac OS 9 version?

There has never been an OS 9 version of SpamSieve—sorry.

8 Contact Information

You can download the latest version of SpamSieve from the SpamSieve Web site[13]. Questions about SpamSieve may be sent to spamsieve@c-command.com[14]. I'm always looking to improve SpamSieve, so please feel free to send any feature requests to that address.

To make sure that you have the latest version of SpamSieve, you may wish to subscribe to the SpamSieve News mailing list[15]. The traffic on this list is very low, only one message per new version of SpamSieve.

9 Purchasing

SpamSieve has a free trial that lasts for 30 days or 20 launches, whichever is longer. To use SpamSieve beyond the demo period, you must purchase it. The price is only $25 (US) and entitles you to free updates and support.

To purchase, choose **Purchase**... from the **SpamSieve** menu. You can use the **Instant Purchase**... button to enter your billing information from within SpamSieve or use the **Web Purchase**... button to enter it at the eSellerate Online Store[16] in your Web browser.

[12] http://homepage.mac.com/fahrenba/dockless/dockless.html
[13] http://www.c-command.com/spamsieve/
[14] mailto:spamsieve@c-command.com
[15] http://www.c-command.com/spamsieve/support.shtml
[16] http://store.eSellerate.net/mt/store

Soon after paying, you'll receive an e-mail with your serial number. Enter the name and serial number from the e-mail into the **Purchase** window and click **Personalize** to personalize your copy of SpamSieve. If you need to re-install SpamSieve, you can simply re-enter your name and serial number and click **Personalize**; there's no need to purchase again.

A license for SpamSieve is good for one person *or* one computer. You can install it on one Mac, and everyone sharing that Mac can use it (on that Mac). Alternatively, you can install it on your desktop Mac and your PowerBook; you can then use it on either machine, provided that no one is using it on the other machine.

SpamSieve uses eSellerate Product Activation[17] to reduce software piracy. This should be completely transparent except that you will need to be connected to the Internet when you first enter SpamSieve's serial number. (Subsequent launches do not require an Internet connection.) eSellerate's privacy policy is as follows:

> eSellerate Product Activation is an anti-piracy technology that publishers can use to protect the software they sell through eSellerate.
>
> During activation, eSellerate looks at the computer's present configuration and uses that data to create a unique hardware identification. The unique hardware identification does not include any personal information, nor does it include any information about the software or documents that reside on the computer. The hardware identification identifies only the computer's configuration, and is used only for activation purposes.
>
> Once software is activated on a computer, minor changes to that computer's configuration will not affect the activation. Major changes to the computer's configuration may require reactivation of the software. Reactivation is at the publisher's discretion.

Purchasing SpamSieve entitles you to a reasonable number of activations. You can activate SpamSieve on your desktop, on your laptop, and on new Macs that you buy in the coming years. If you run out of activations, e-mail me and I'll most likely give you more. The goal is to make things as easy as possible for owners of SpamSieve.

10 Version History

2.0—September 10, 2003

- SpamSieve now extracts *a lot* more information from each message. This makes it much more accurate and also makes it learn faster.

- Now integrates with Eudora 6 (Sponsored or Paid) via a plug-in. It can now process every incoming Eudora message and can be trained using the **Junk** and **Not Junk** commands in Eudora's **Message** menu.

[17] http://www.esellerate.net/papolicy.asp

- SpamSieve now has a blocklist and a whitelist. These are automatically maintained based on the senders of messages that SpamSieve is trained with. The blocklist makes sure that all messages from known spammers are caught and speeds processing for these messages. The whitelist lets you be sure that certain messages will never be marked as spam; this was possible before, but now you don't have to clutter your address book with addresses from online retailers, etc.
- You can now control how conservative or aggressive SpamSieve is at catching spam.
- SpamSieve can now play a sound or bounce its Dock icon after a batch of non-spam messages has arrived. This is meant to replace your e-mail client's new mail notification, which you don't want going off if all the new messages are spam.
- Shows the number of new good messages in the Dock icon.
- Now parses HTML so that it can better extract relevant information from HTML messages, and also handle various HTML-based tricks that spammers use to fool filters.
- New method of calculating word probabilities makes SpamSieve better at discerning which words in the message are important.
- Includes a corpus of seed spam, to jump-start spam recognition for users who do not have many saved spam messages.
- The corpus is now stored in databases rather than in a property list. This makes it launch faster and use much less memory, as the corpus doesn't have to be all in RAM at the same time.
- The statistics file format (for History.db) has changed in order to enable performance improvements and more statistical displays in future versions.
- Handles more types of plain text obfuscations, and is much faster at undoing them.
- Added option for the address book whitelist to only use other people's addresses, so that spam messages from your own address don't match the whitelist.
- Can mark all messages with Habeas headers as good.
- Can mark all messages with some variant of "ADV" at the start of the subject as spam.
- Can mark all base64-encoded HTML messages as spam.
- New probability combiner increases accuracy.
- Uses stop words to speed processing and reduce false negatives.
- When filtering a message, considers the number of occurrences of the words, not just which words are present.
- Can import messages from mbox files.

- Can import the corpus from and export it to an XML property list (the same format used by 1.x).
- SpamSieve can now check for updated versions of itself.
- Added crash reporter.
- Added Dock menu containing frequently used commands.
- The entries in the log are more detailed.
- The corpus now stores the date at which each word was last accessed.
- Fixed bug where storing statistics would fail on systems that didn't know about GMT.
- Fixed bug where SpamSieve could throw away long runs of HTML thinking they were attachments.
- Added button for opening the Mac OS X Address Book from inside SpamSieve.
- The **Statistics** window now has a contextual menu item for copying the displayed information.
- SpamSieve no longer wastes cycles updating the **Statistics** window after it's been closed.
- The **Statistics** window is smarter about updating only the portions that could have changed.
- No longer shows Good Words and Spam Words stats.
- Logging has less overhead.
- Updates the history asynchronously, resulting in faster message processing.
- Checks for mistakes in a background thread.
- False negatives are now written to disk in a background thread.
- Re-arranged the **Corpus** window.
- Pruning the corpus now works by access date rather than by word counts. Of course, you can manually prune the old way by sorting the **Corpus** window by **Total**.
- Updated to SQLite 2.8.6 and tuned it for speed.
- Updated to PCRE 4.3.
- Updated to eSellerate 3.5, which should fix crashes some people saw after registering on 10.2.6.

- Now looks at headers of subparts of messages from Mailsmith.
- Time-consuming operations now either have a progress bar or a progress spinner.
- Better at extracting malformed e-mail addresses from headers.
- Copying rows from the **Corpus** window to the clipboard now uses the order of the columns in the window rather than the default column order.
- Fixed regression where the Entourage scripts no longer created the **Spam** folder if it didn't exist.
- Fixed potential crash with regex replacements at the end of a string.
- History.db and the corpus can now be aliases.
- Automatically trims carriage returns and other illegal characters when you paste in your name and serial number.
- Now saves the name and serial number to disk as soon as they're entered.
- The **Spam** folder in Entourage no longer has to be top-level.
- Entourage can mark good messages as unread.
- Type-selecting in table views is quicker.
- No longer nags constantly when unregistered.
- Fixed bug where it could *look* as though SpamSieve had hung if it started up in the background with an empty corpus.

1.3.1—June 18, 2003

- Added direct integration with Mailsmith 2.0 and later. Enabling SpamSieve is as easy as clicking a checkbox. You can train SpamSieve directly from Mailsmith's Message menu. Bare Bones Software has seamlessly integrated it with Mailsmith's powerful filtering system, and Mailsmith knows not to bounce its Dock icon after receiving a batch of messages that are all spam.
- Fixed crashing bug triggered by incorrectly encoded headers.
- Regex substitutions are faster and much more memory efficient.
- When adding spam messages to the corpus, the default is now for SpamSieve to move them to the **Spam** folder.
- The PowerMail **Move If Spam** script now changes the color of spam messages.

- The Emailer scripts now pass text and HTML attachments on to SpamSieve for analysis.
- Added instructions for using the Entourage and PowerMail address books as whitelists.
- Compacted the ED frameworks to reduce application size and memory use.
- Disabled SQLite's file locking so that SpamSieve's data folder can now be located on an AppleShare volume.
- Caches the Address Book to speed whitelist lookups 100 fold.
- The statistics database is faster due to an updated version of SQLite.
- Fixed bad pointer in header extractor that caused unnecessary fallbacks to the secondary parser (reducing accuracy).
- Added experimental support for moving Entourage messages on IMAP accounts to the (local) **Spam** folder. This can be enabled by editing the scripts with Script Editor.
- For clarity, the names of AppleScripts that ask SpamSieve to predict the category of a message now contain the word "if."
- Added lots of minor clarifications to the documentation.

1.3—February 11, 2003

- More resilient to spammers' tricks for obfuscating words.
- Can use e-mail addresses in the system Address Book as a whitelist. Messages sent from those addresses will never be marked as spam.
- Greatly reduced overall memory usage as well as launch and quit times.
- Can save false negatives to disk for later reporting to SpamSieve's developer.
- You can edit the spam and good counts associated with a word, remove selected words from the corpus, and reset the corpus entirely.
- Type-ahead navigation in the **Corpus** window. Type the first few letters of a word or number to select it (and scroll to it).
- You can hide statistics from before a set date, to better see the current accuracy and spam reception rate.
- Improvements to the **Corpus** window: Shows all words rather than only those considered statistically significant. Re-sorting by numeric columns is twice as fast. You can copy the selected rows to the clipboard or drag them to another application. The selection is preserved when you change the sort column, you can sort in descending order, and the sorted column is remembered between launches. The Home and End keys work.

- The **Prune Corpus** command now tells you how many words it would remove and asks for confirmation.
- The statistics tracking is smarter about handling duplicate messages.
- The statistics have tooltips explaining what they mean, and you can copy all of the statistics to the clipboard at once.
- Improved accuracy tracking of PowerMail and Emailer messages.
- Eudora Integration: Can mark spam messages as read and/or mark them for removal from the server.
- Expanded the AppleScript dictionary, to enable better integration with mail and news clients.
- Entourage Integration: Creates Junk category if there isn't one, and can mark spam messages as read.
- Mailsmith Integration: The adding scripts now set the appropriate message properties.
- Better parsing of messages with illegal characters in the headers.
- SpamSieve's Info.plist file contains an `LSUIElement` entry. Change the `0` to a `1` to hide the application's Dock icon. (You'll need to change it back to access the preferences.)
- The message count display has moved from the **Corpus** window to the **Statistics** window.
- Better error message when the corpus couldn't be saved.
- Added tooltips to preferences.
- The registration window gives better feedback when you personalize.
- Better recovery from errors in the corpus file.
- The secondary parser is better at handling DOS linebreaks.

1.2.2—November 20, 2002

- Fixed bug in the PowerMail **Add Good** script.
- Added uninstaller for Eudora users.
- Better handling of errors while adding messages to the corpus.
- Removed bloat from the Entourage **Mark Spam** script.
- The application icon now has an alpha channel, so it doesn't appear with a white halo when viewed on a colored background.
- Minor changes to the manual.

1.2.1—November 18, 2002

- Modified Info.plist to work around a bug in Mac OS X 10.1 that could cause the Finder to crash when launching SpamSieve.

1.2—November 18, 2002

- Added support for Emailer 2.0v3 and Eudora (5.2 and later).
- Decodes base64 and quoted-printable text parts, thus finding words that spammers try to hide from anti-spam software.
- Decodes subjects that use different character sets (e.g. big5).
- Adds special tokens for MIME entities such as part boundaries and uninterpretable message parts.
- Keeps track of the messages added to the corpus, and can optionally prevent you from adding the same message more than once (biasing the counts). Thus, you no longer have to remember which messages you've already added.
- You can now "undo" adds to the corpus, e.g. if you added a message as good when you meant to add it as spam.
- Can now add messages to the corpus as they are filtered, so after the initial training you only have to add messages when SpamSieve makes a mistake.
- When filtering a message, SpamSieve can optionally check whether the message is in the corpus. If it is, SpamSieve looks up the answer rather than trying to predict. One use of this feature is that if SpamSieve makes a mistake, you can **Add Spam** and then **Label/Move If Spam** and be sure that the message will be labelled/moved.
- Keeps a log of additions to the corpus, filtering results, and errors.
- Mailsmith: If SpamSieve thinks a message is spam, it sets the **deleted** property of the message to true; otherwise it sets the **flagged** property of the message to true. Therefore, if SpamSieve has classified the message then exactly one of the properties will be true, and if it hasn't they'll both be false. (Normally, neither of these message properties is used by Mailsmith itself.)
- Entourage and PowerMail: If you tell SpamSieve to move spam messages to a spam folder and the spam folder doesn't exist, the script will create the spam folder for you.
- When you add spam messages to the corpus, can optionally move them to a **Spam** folder.
- Added status indicators in the Dock icon (like Norton DiskLight).
- The spam probability of unknown words is now 0.4 instead of 0.2.

- The **Corpus** window uses less memory and sorts much faster.
- Accuracy tracking is faster and uses less memory and disk space.
- Fixed bug where accuracy tracking didn't work for some Mailsmith messages with multiple parts.
- Improved the manual's instructions for e-mail client integration.
- Compiled with GCC 3 for greater speed.
- Uses the latest version of the eSellerate SDK, which eliminates a crash at startup under certain circumstances.
- No longer shows the "Upgrading From 1.0" message when starting with a blank corpus.

1.1—September 19, 2002

- E-Mail Client Integration
  - Added support for PowerMail.
  - Added instructions and an AppleScript for making Mailsmith download and filter mail faster.
  - Added an AppleScript for Entourage that moves spam into a Junk folder.
- Performance
  - Launches about 60% faster than 1.0.
  - You can now prune the corpus to remove words that are taking up memory without contributing to spam recognition. This can also dramatically decrease SpamSieve's launch time.
  - Recalculating spam probabilities is about 10% faster and uses less memory.
  - Quitting is faster because SpamSieve now writes corpus changes to disk during idle time.
  - Saving the corpus is slightly faster.
- Displays statistics about the number of messages filtered, SpamSieve's accuracy, and the types of words in the corpus.
- SpamAssassin's X-Spam-Status headers are now treated as single words. This means that if SpamAssassin is running on your mail server, SpamSieve will learn to respect (or ignore) its judgement.
- Does a better job of ignoring e-mail attachments, thus reducing corpus bloat.
- Installs the eSellerate Engine if it's not present, thus enabling "Instant Registration" for more users.

- Asking SpamSieve to categorize a message now forces an update of all the word probabilities. Previously, the update only happened during idle time.
- Highlights the sorted column in the **Corpus** window. The columns themselves have shorter names. There's a new "Total" column. Auto-resizing of the columns works better. You can now manually resize any column, and manual resizings and reorderings are saved between launches.
- Shows fatal errors as alert panels rather than just printing them on the console.
- The Corpus.plist data file is now sorted by word. This makes it easier to examine the corpus manually, and to compare it to other users' corpuses.

1.0—September 10, 2002

- First public release.

11 Legal Stuff

SpamSieve License Agreement

PCRE

Regular expression support is provided by the PCRE[19] library package, which is open source software, written by Philip Hazel, and copyright by the University of Cambridge, England.

[18] mailto:mjt@c-command.com
[19] http://www.pcre.org

EDCommon

EDCommon[20] is Copyright © 1999–2002 by Erik Doernenburg. All rights reserved.

Permission to use, copy, modify and distribute this software and its documentation is hereby granted, provided that both the copyright notice and this permission notice appear in all copies of the software, derivative works or modified versions, and any portions thereof, and that both notices appear in supporting documentation, and that credit is given to Erik Doernenburg in all documents and publicity pertaining to direct or indirect use of this code or its derivatives.

EDMessage

EDMessage[21] is Copyright © 2000–2002 by Erik Doernenburg and Axel Katerbau. All rights reserved.

Permission to use, copy, modify and distribute this software and its documentation is hereby granted, provided that both the copyright notice and this permission notice appear in all copies of the software, derivative works or modified versions, and any portions thereof, and that both notices appear in supporting documentation, and that credit is given to Erik Doernenburg in all documents and publicity pertaining to direct or indirect use of this code or its derivatives.

A Mailsmith Extras

Mailsmith 2.0 and later integrate directly with SpamSieve. This is described in Section 4.4.1.

SpamSieve also includes a `Mailsmith Extras` folder, which contains AppleScripts for use with Mailsmith. These make it possible for scripters to further customize and automate the labelling and marking of messages in Mailsmith. If you are using Mailsmith 2.0 and do not write AppleScripts, you can ignore the `Mailsmith Extras` folder.

You can add the AppleScripts to Mailsmith's **Scripts** menu by copying them to the `Scripts` folder inside the `Mailsmith Support` folder. A filter such as the following may be used to change the labels of incoming spam messages.

[20] http://www.mulle-kybernetik.com/software/EDFrameworks/download.html#EDCommon
[21] http://www.mulle-kybernetik.com/software/EDFrameworks/download.html#EDMessage

This filter will pass all messages along to SpamSieve for analysis. It will set the **Is Spam** or **Is Not Spam** property of the message, and change the label of the message if it is spam. This is roughly equivalent to enabling SpamSieve in Mailsmith's preferences, but because it uses AppleScript it is more customizable. Additionally, you can change the filter criteria to pass only select messages along to SpamSieve.

For best results, use either Mailsmith's direct integration with SpamSieve or AppleScripts like those in the `Mailsmith Extras` folder. Do not mix and match them.

B Using SpamSieve With Eudora 5.2

B.1 Installing

Using SpamSieve with Eudora 6 (Section 4.3.1) is highly recommended. However, SpamSieve can also work with Eudora 5.2, and some Eudora 6 users may prefer the configuration described here because it is more customizable.

Copy the `SpamSieve Eudora Helper` file to the `Applications` folder of your hard disk. You will need to launch this applet the first time you use SpamSieve with Eudora. Thereafter, Eudora should launch the applet automatically. Do not launch the copy on the disk image, as that will cause Eudora to try to find it in the wrong place.

You can copy the `Uninstall Eudora Helper` file to your hard disk, or just leave it on the disk image. Run this applet if you no longer want to use SpamSieve with Eudora.

Create a mailbox in Eudora called **Spam** that is at the same level as the **In** mailbox. When you receive new spam messages, SpamSieve will move them to the **Spam** mailbox. It will also mark good messages by setting their priority to lowest (indicated by two downward pointing carets) and mark spam messages by setting their status to transfer error (indicated by a red "X").

B.2 Training SpamSieve

To train SpamSieve with spam messages, select one or more of them. Choose the **Filter Messages** command in Eudora's **Special** menu. Then double-click **Add Spam**.

To train SpamSieve with good messages, select one or more of them. Choose the **Filter Messages** command from Eudora's **Special** menu. Then double-click **Add Good**.

B.3 Manually Processing Messages

To manually ask SpamSieve to mark or move messages that it thinks are spam, select one or more of them. Choose the **Filter Messages** command in Eudora's **Special** menu. Then double-click **Filter**.

B.4 Setting Options

By configuring the `SpamSieve Eudora Helper` applet, you can tell SpamSieve to process your good messages and spam messages in other ways. First, quit the applet. Then open it using the Script Editor program in the `AppleScript` folder of your `Applications` folder. The top of the script contains the following lines:

```
property moveToSpamFolder     : true  -- moves spams to a "Spam" mailbox
property markSpamMessages     : true  -- marks spams with red x
property markSpamMessagesRead : false -- marks spams as "already read"
property labelSpamMessages    : false -- colors spam messages brown
property markGoodMessages     : true  -- marks good messages with carets
property labelGoodMessages    : false -- colors good messages green
property removeSpamMessagesFromServer : false
```

You can change a `false` to `true` or a `true` to `false` to set the options the way you want. For instance, to have SpamSieve not move spam messages into a separate mailbox, change the `true` in the first line to `false`. When you are finished making changes, choose **Save** in Script Editor's **File** menu, close the window, and then re-launch the SpamSieve Eudora Helper.

B.5 Eudora Limitations

The following limitations are due to problems with Eudora's "notification" interface. Because of these limitations it is recommended that you use Eudora 6 and the `SpamSieve Eudora Plug-In`, as described in Section 4.3.1. The plug-in avoids these limitations.

- Eudora gives messages to SpamSieve *after* all the other filters have run. It is not possible to change this ordering.
- SpamSieve cannot filter messages that are moved by other filters. For instance, if you have a filter that moves incoming messages from Steve Jobs to a separate mailbox, SpamSieve will not mark any of those messages as spam, even if a spammer pretends to be Jobs. This limitation applies to both automatic filtering of incoming mail and manual filtering of selected messages.
- Sometimes the wrong message is marked. That is, SpamSieve may decide that message A is spam and ask Eudora to mark it with a red "X"; in rare circumstances, Eudora will instead mark some other message B with the "X." You can tell if this has happened by comparing SpamSieve's log to the way the messages are marked in Eudora. This problem seems to occur when the **In** mailbox sorted.
- Sometimes SpamSieve never sees a message that should have been filtered. You can tell if this has happened by the absence of that message in the log. It may help to remove any "notify user" filter action that you have set up.
- Sometimes SpamSieve determines that a message is good or spam, but Eudora does not mark it all. You can tell if this has happened by comparing SpamSieve's log to the way the messages are marked in Eudora.
- SpamSieve cannot add or filter messages that are stored in the Trash mailbox or in mailbox files outside the `Mail Folder` folder in the `Eudora Folder`. Note that this includes all IMAP messages. To access these messages, first move them to a non-trash mailbox file that is stored inside the `Mail Folder` folder.
- If you manually apply filters while Eudora is in the process of downloading mail, Eudora will show the SpamSieve dialog box twice. If this happens, just choose **Skip** the second time.
- Sometimes Eudora erroneously shows the SpamSieve dialog when you check for new mail.
- Sometimes after a long delay in talking to the mail server, Eudora stops notifying SpamSieve when it receives new messages. You can work around this by quitting and re-launching the SpamSieve Eudora Helper.