

# SpamSieve 1.3 Manual

Michael Tsai

February 10, 2003



# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	What Is SpamSieve? . . . . .	3
1.2	Identifying Spam . . . . .	3
1.3	Filtering It Out . . . . .	4
<b>2</b>	<b>Requirements</b>	<b>4</b>
<b>3</b>	<b>Installation</b>	<b>5</b>
3.1	The SpamSieve Application . . . . .	5
3.2	Mailer Integration Scripts . . . . .	5
3.3	Entourage Integration Scripts . . . . .	6
3.4	Eudora Integration Applet . . . . .	6
3.5	Mailsmith Integration Scripts . . . . .	6
3.6	PowerMail Integration Scripts . . . . .	7
<b>4</b>	<b>Training SpamSieve to Recognize Your Spam</b>	<b>7</b>
<b>5</b>	<b>Using SpamSieve With Your E-Mail Client</b>	<b>8</b>
5.1	General Filtering Advice . . . . .	8
5.2	Mailer . . . . .	8
5.2.1	Basic Usage . . . . .	8
5.2.2	Coloring Predicted Spam Messages . . . . .	9
5.2.3	Moving Predicted Spam Messages . . . . .	9
5.3	Entourage . . . . .	10
5.3.1	Basic Usage . . . . .	10
5.3.2	Marking Predicted Spam Messages . . . . .	10
5.3.3	Moving Predicted Spam Messages . . . . .	11
5.3.4	IMAP Accounts . . . . .	12
5.3.5	Advanced Rules . . . . .	12
5.4	Eudora . . . . .	13
5.4.1	Basic Usage . . . . .	13
5.4.2	Marking Predicted Good and Spam Messages . . . . .	13
5.4.3	Moving Predicted Spam Messages . . . . .	14
5.4.4	Eudora Limitations . . . . .	14
5.5	Mailsmith . . . . .	15
5.5.1	Basic Usage . . . . .	15
5.5.2	Labelling Predicted Spam Messages . . . . .	15
5.5.3	Moving Predicted Spam Messages . . . . .	16
5.5.4	Advanced Filtering . . . . .	16
5.5.5	Improving Performance on Mac OS X 10.2 and 10.2.1 . . . . .	16
5.6	PowerMail . . . . .	17
5.6.1	Basic Usage . . . . .	17
5.6.2	Labelling Predicted Spam Messages . . . . .	17

5.6.3	Moving Predicted Spam Messages . . . . .	17
5.7	Customizing the AppleScripts . . . . .	18
5.7.1	Compatibility Notes . . . . .	18
5.7.2	Moving Spam Messages After Adding Them . . . . .	18
5.7.3	Integrating With Other Applications . . . . .	18
<b>6</b>	<b>The Filter Menu</b>	<b>18</b>
6.1	Show Corpus . . . . .	18
6.2	Remove Word... . . . .	19
6.3	Prune Corpus... . . . .	19
6.4	Show Statistics . . . . .	19
6.5	Open Log . . . . .	20
<b>7</b>	<b>Preferences</b>	<b>20</b>
7.1	Allow duplicate good/spam messages . . . . .	20
7.2	First check if message is in corpus . . . . .	21
7.3	Add predicted good/spam messages to corpus . . . . .	21
7.4	Use system Address Book as a whitelist . . . . .	21
7.5	Save false negatives to disk . . . . .	22
<b>8</b>	<b>Contact Information</b>	<b>22</b>
<b>9</b>	<b>Registering</b>	<b>22</b>
<b>10</b>	<b>Notes and Limitations</b>	<b>23</b>
<b>11</b>	<b>Version History</b>	<b>23</b>
<b>12</b>	<b>Legal Stuff</b>	<b>27</b>

# 1 Introduction

## 1.1 What Is SpamSieve?

SpamSieve is an application that integrates with e-mail clients to filter 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. People have turned to a wide variety of software solutions to help regain control over their inboxes. This software differs primarily along two dimensions: how it identifies spam messages; and how it reduces their burden on the user.

## 1.2 Identifying Spam

*Simple rules* that you create in your e-mail client generally look for common spam words. If a message’s subject contains “fix your credit” or its body contains “free porn” it’s probably spam. You can look at the spam you receive and recognize some patterns. Unfortunately, spam is constantly evolving, so you will have to keep working on your rules. Eventually you get to the point where you don’t know what else to add, afraid that adding more words or phrases to your “blacklist” would start filtering out legitimate mail. You can somewhat get around this by creating a “whitelist” that accepts every message from the people you regularly correspond with. However, this doesn’t help for messages received from new people, and it’s not uncommon for spam messages to contain a forged return address of someone at your own company.

Commercial anti-spam software often combines rules with the notion of a *score*. Rules can look for patterns that make a message more or less likely to be spam. Each rule either increases or decreases the message’s spam score. If the score is above a certain threshold, the message is considered to be spam. The flexible nature of the score means that this approach is often an improvement over simple rules. However, the user typically has little or no control over how the spam score is calculated. If a spam message gets through, the user has no recourse but to hope that the next update corrects the problem. Another problem is that spammers can buy this software, too. They can tailor their spam to get through the rules.

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](http://www.paulgraham.com/spam.html)<sup>1</sup> and the papers it references. Bayesian analysis combines the good properties of the above two approaches and adds some of its own. 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*

---

<sup>1</sup><http://www.paulgraham.com/spam.html>

*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 Filtering It Out

Anti-spam software that runs on mail servers filters out spam before you ever see it. This means that unless the filter is perfect, either some spam messages will get through or a few legitimate messages will not. In the first case, you may want additional, client-side, anti-spam software. The second case troubles many people so much that they prefer that the server do no filtering at all.

Other client-side anti-spam software connects to your mail server to delete spam messages before your e-mail client can download them. This works similarly to above, except that 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 and also download good messages 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.

E-mail clients like Entourage have their own spam filters built-in. This is convenient and makes it easy to manually scan for false positives, but there is typically little you can do when the filter makes mistakes.

Clearly, the best solution is user-configurable anti-spam software that works directly with your e-mail client. Apple has added exactly this to its Mail.app program in Mac OS X 10.2. By most accounts this works great...unless you prefer a client other than Mail.app. This is where SpamSieve comes in. It brings powerful spam filtering to other popular e-mail clients. SpamSieve works with any number of mail accounts and supports whatever account types (POP, IMAP, Hotmail, AOL) are supported by your e-mail client.

## 2 Requirements

SpamSieve has been developed and tested on Mac OS X 10.1.5 and 10.2.3. I do not have the resources to test it on older systems, although I suspect it will work fine on Mac OS X 10.1 or later. SpamSieve is designed to work with the following e-mail clients:

- [Emailer 2.0v3<sup>2</sup>](#), previously available from Claris

---

<sup>2</sup><http://www.fogcity.com>

- [Entourage 9.0.1](#)<sup>3</sup> and later from Microsoft
- [Eudora 5.2](#)<sup>4</sup> and later from Qualcomm
- [Mailsmith 1.5](#)<sup>5</sup> and later from Bare Bones Software
- [PowerMail 4.0](#)<sup>6</sup> and later from CTM Development

SpamSieve is best used with Mailsmith, due to Mailsmith's flexible filter criteria and high level of support for AppleScript filter actions.

## 3 Installation

### 3.1 The SpamSieve Application

Copy the SpamSieve application from the `SpamSieve-1.3.dmg` disk image to your hard disk, e.g. into `/Applications`. 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.

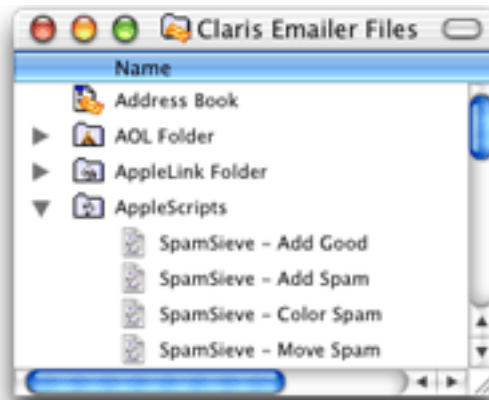
Find the folder on the disk image that pertains to your e-mail client (e.g. **For Entourage Users**). The AppleScripts in this folder allow you to interact with SpamSieve from within your e-mail client. The names of the contained folders tell you where to copy their contents.

*Be sure that you replace any existing scripts from a previous version of SpamSieve. You may need to reconfigure the mail rules/filters in your e-mail client so that they reference the new scripts.*

Follow the instructions below that correspond to your e-mail client.

### 3.2 Mailer Integration Scripts

Make sure that you have installed the AppleScripts as shown here:



<sup>3</sup><http://www.microsoft.com/mac/entourage/default.asp?navindex=s4>

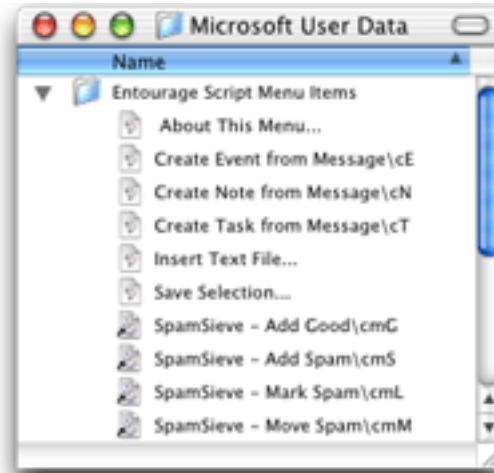
<sup>4</sup><http://www.eudora.com/mac>

<sup>5</sup><http://www.barebones.com/products/mailsmith.html>

<sup>6</sup><http://www.ctmdev.com/powermail4.shtml>

### 3.3 Entourage Integration Scripts

Disable Entourage's Junk Mail Filter. Then make sure that you have installed the AppleScripts as shown here:



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

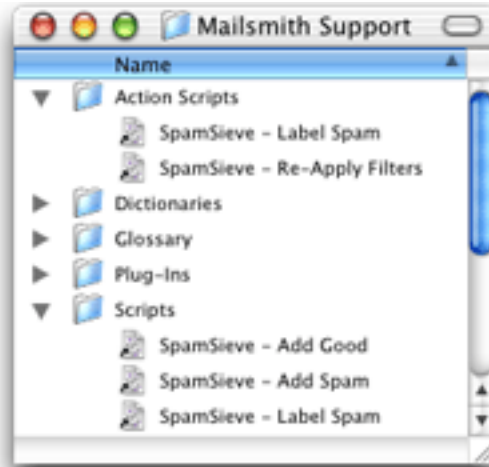
### 3.4 Eudora Integration Applet

Make sure that you have copied the **SpamSieve Eudora Helper** file to your hard disk. You will need to launch this applet any time you want to use SpamSieve from Eudora. If you want, you can use System Preferences to make it a login item. 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.

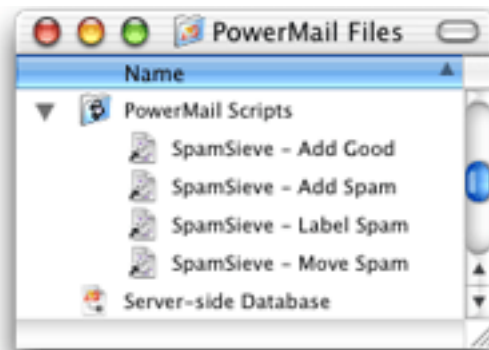
### 3.5 Mailsmith Integration Scripts

Make sure that you have installed the AppleScripts as shown here:



### 3.6 PowerMail Integration Scripts

Make sure that you have installed the AppleScripts as shown here:



## 4 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 5. 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 very conservative about identifying spam, leading to more false negatives. Most users find that filling the corpus with a ratio of 2:1 to 4:1 spam messages to good messages (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, accuracy will deteriorate over time.*

Entourage users: make sure that you correct SpamSieve’s mistakes by running the AppleScripts in your Scripts menu—do not click the underlined blue text in Entourage to indicate that a message is not spam.

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 add new spam messages to its corpus as you receive them. It does not automatically add new 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 add messages to its corpus to correct mistakes.

## 5 Using SpamSieve With Your E-Mail Client

### 5.1 General Filtering Advice

Most of the e-mail clients that SpamSieve supports let you control the order in which their rules (a.k.a. filters or mail actions) 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’d 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.

### 5.2 Mailer

#### 5.2.1 Basic Usage

- To add spam messages to SpamSieve, select one or more of them and run the **SpamSieve - Add Spam** script from the **Scripts** menu.
- To add good messages to SpamSieve, select one or more of them and run the **SpamSieve - Add Good** script from the **Scripts** menu.

- SpamSieve will mark or move new spam messages that you receive, depending on how you configure EMailer below.
- To manually color or move spam messages, select them and run the **SpamSieve - Color Spam** or the **SpamSieve - Move Spam** script from the **Scripts** menu.

### 5.2.2 Coloring Predicted Spam Messages

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



### 5.2.3 Moving Predicted Spam Messages

If, instead, you want SpamSieve to move suspected spam messages to a **Spam** folder, set up a mail action in EMailer that looks like this:



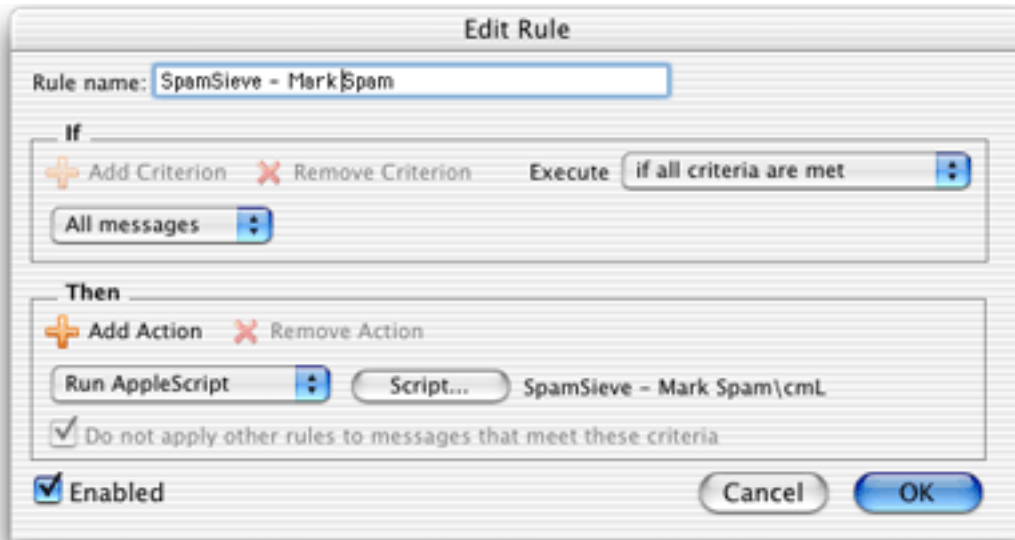
## 5.3 Entourage

### 5.3.1 Basic Usage

- To add spam messages to SpamSieve, select one or more of them and run the **SpamSieve - Add Spam** script from the **Scripts** menu.
- To add good messages to SpamSieve, select one or more of them and run the **SpamSieve - Add Good** script from the **Scripts** menu.
- SpamSieve will mark or move new spam messages that you receive, depending on how you configure Entourage below.
- To manually mark or move spam messages, select them and run the **SpamSieve - Mark Spam** or the **SpamSieve - Move Spam** script from the **Scripts** menu.

### 5.3.2 Marking Predicted Spam Messages

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

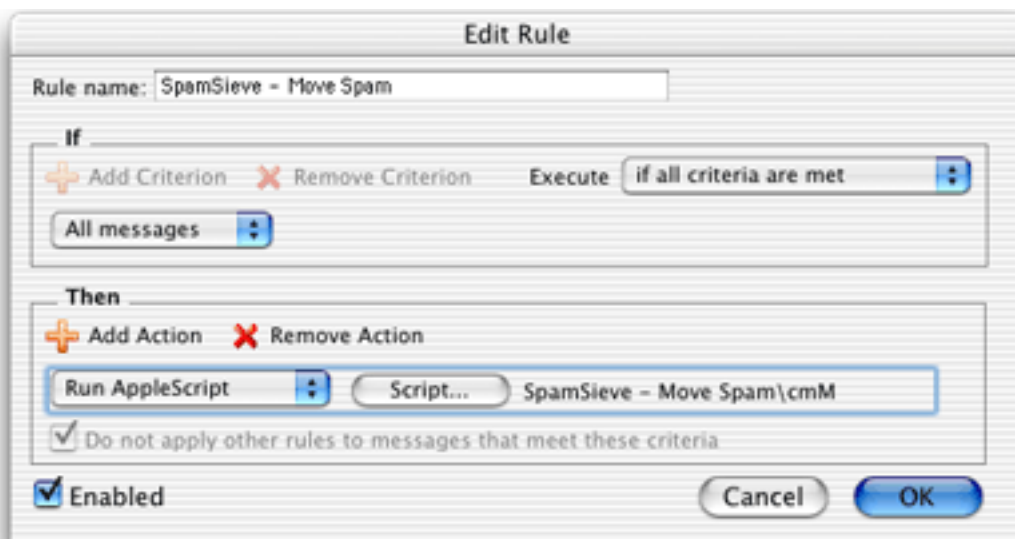


Be sure to create the kind of rule that is appropriate for your account type, e.g. a POP rule for a POP account or an IMAP rule for an IMAP account. If you have more than one kind of account, you will need to create a rule for each account type.

It is important that you create the rules 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.

### 5.3.3 Moving Predicted Spam Messages

If, instead, you want SpamSieve to move suspected spam messages to a **Spam** folder, set up the mail rule to use the **SpamSieve - Move Spam** script instead.



### 5.3.4 IMAP Accounts

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

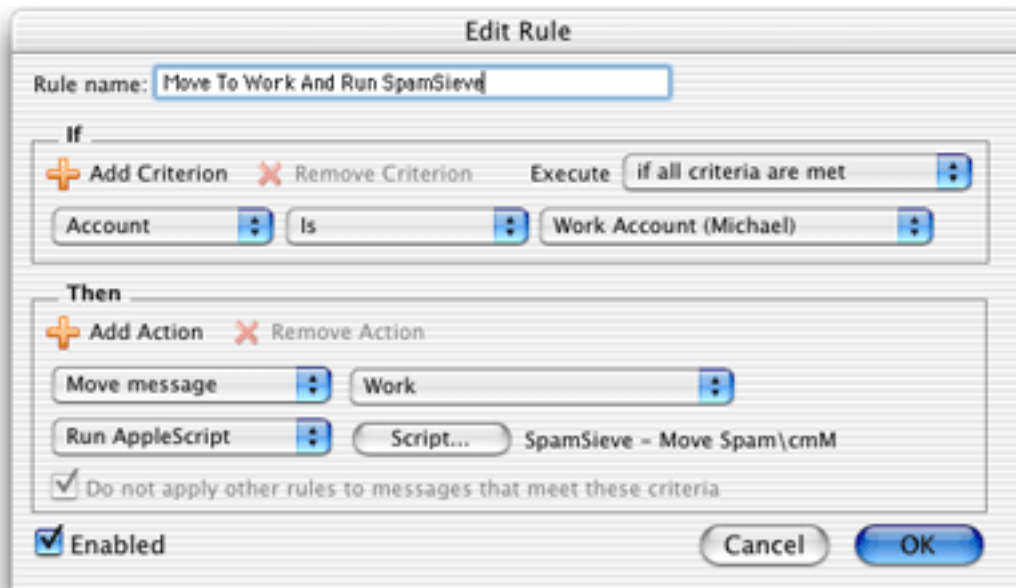
Yet another alternative is to use a custom view. Make a custom view of your IMAP account that looks for messages that are not junk. In this way, you can view your good messages without being distracted by spams.

### 5.3.5 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. 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 SpamSieve 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.

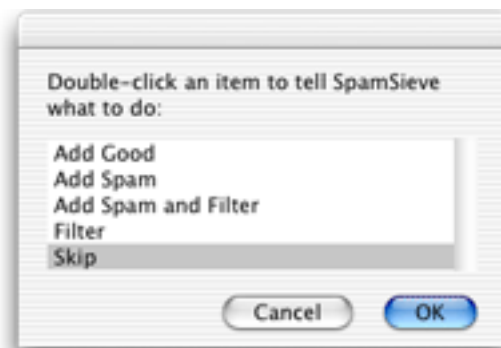
You can add the SpamSieve AppleScript action to every rule that moves messages and also to a “catch-all” rule that applies to messages that aren’t moved. Then SpamSieve will be able to filter all the messages that you receive.

Please contact [spamsieve@c-command.com](mailto:spamsieve@c-command.com)<sup>7</sup> if you have trouble setting up Entourage to filter messages the way you want.

## 5.4 Eudora

### 5.4.1 Basic Usage

- To add spam messages to SpamSieve, select one or more of them. Choose the **Filter Messages** command in the **Special** menu. Then double-click **Add Spam**.



- To add good messages to SpamSieve, select one or more of them. Choose the **Filter Messages** command from the **Special** menu. Then double-click **Add Good**.
- SpamSieve will mark or move new spam messages that you receive, depending on how you configure the SpamSieve Eudora Helper below.
- To manually mark or move spam messages, select one or more of them. Choose the **Filter Messages** command in the **Special** menu. Then double-click **Filter**.

### 5.4.2 Marking Predicted Good and Spam Messages

By default, SpamSieve will 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”). You can sort Eudora’s **In** mailbox by message status (the first column) to separate the spam messages from good ones.

---

<sup>7</sup><mailto:spamsieve@c-command.com>

### 5.4.3 Moving Predicted Spam Messages

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 Script Editor. The top of the script contains the following lines:

```
property moveToSpamFolder      : false -- 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 move spam messages into a separate mailbox, change the `false` in the first line to `true`.

### 5.4.4 Eudora Limitations

The following limitations are due to problems with Eudora's "notification" interface. They are current as of Eudora 5.2. Qualcomm has been notified of these issues and hopefully will resolve them in a future version of Eudora.

- 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 is sorted by date and the new message is dated earlier than the newest message in the mailbox. You can somewhat work around this by moving messages out of the **In** mailbox after reading them.
- 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.

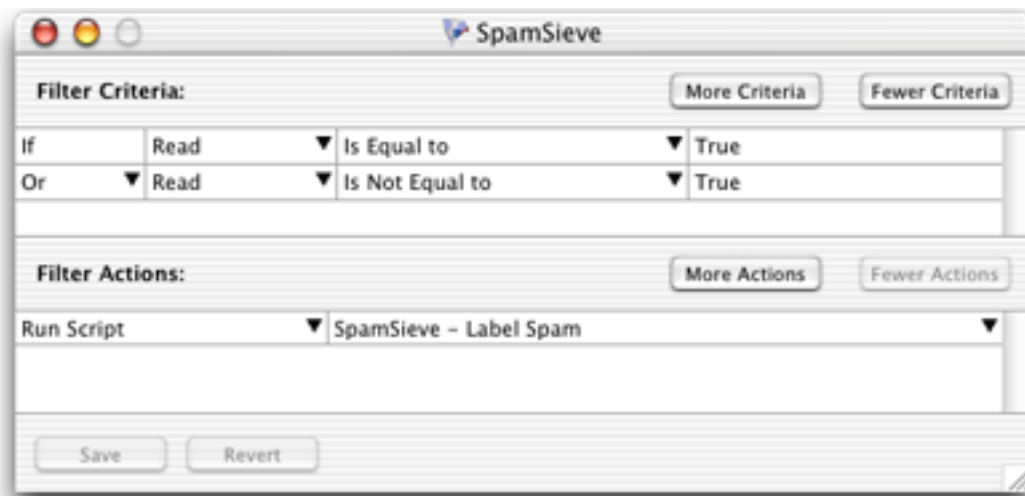
## 5.5 Mailsmith

### 5.5.1 Basic Usage

- To add spam messages to SpamSieve, select one or more of them and run the **SpamSieve - Add Spam** script from the **Scripts** menu.
- To add good messages to SpamSieve, select one or more of them and run the **SpamSieve - Add Good** script from the **Scripts** menu.
- SpamSieve will label or move new spam messages that you receive, depending on how you configure Mailsmith below.
- To manually label spam messages, select them and run the **SpamSieve - Label Spam** script from the **Scripts** menu.

### 5.5.2 Labelling Predicted Spam Messages

Create a filter that looks like this:

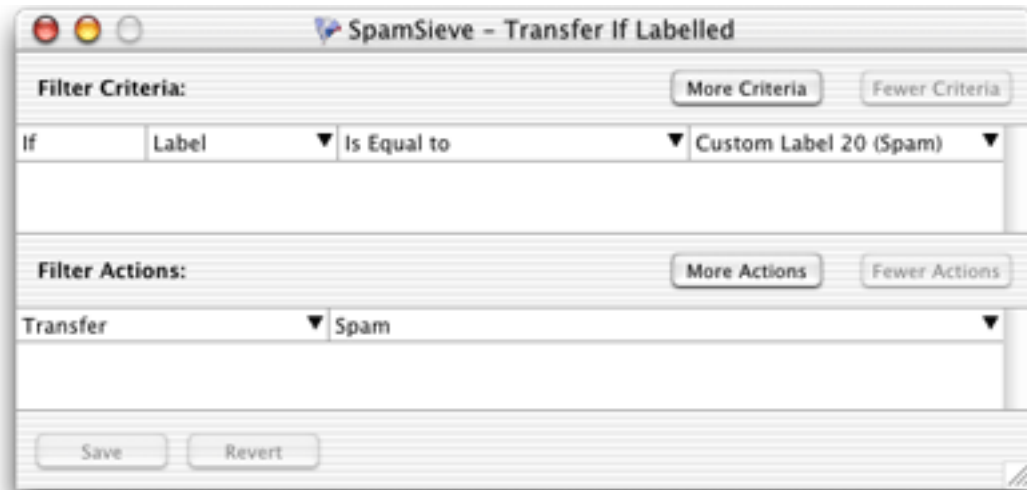


Attach it to the **(incoming mail)** mailbox.



### 5.5.3 Moving Predicted Spam Messages

In addition to the above filter, create a filter such as the following to transfer spam messages to a particular mailbox:



Attach it to the (**incoming mail**) mailbox, and make sure that this filter is applied after the one that uses the **SpamSieve - Label Spam** script.

### 5.5.4 Advanced Filtering

When SpamSieve thinks that a message is spam it will set its label to Custom Label 20. You can then create additional filters that move or otherwise process messages with this label. For more advanced filtering, SpamSieve also does the following:

- If SpamSieve thinks a message is spam, it sets the **deleted** property of the message to true. This property is not visible in Mailsmith's mail lists, but it is available from filter criteria and advanced queries.
- If SpamSieve thinks a message is good, it sets the **flagged** property of the message to true. This property is not visible in Mailsmith's mail lists, but it is available from filter criteria and advanced queries.

If SpamSieve has not examined a message, both of these properties will be false. This allows you to create filters and queries that differentiate between messages that SpamSieve has not seen and those which SpamSieve has examined and found to be good.

### 5.5.5 Improving Performance on Mac OS X 10.2 and 10.2.1

Due to problems in these versions of the operating system, using Mailsmith and SpamSieve together will result in slowness when downloading and filtering mail. For this reason, it is recommended that you use either Mac OS X 10.1.5 or 10.2.3.

If you need to use 10.2 or 10.2.1, there are a few steps you can take to improve performance. First, go to Mailsmith's **Filtering** preferences and uncheck **Apply Filters while Downloading**. Next, go to Mailsmith's **Notification** preferences. Check the **Run Script** checkbox and from the popup menu select the **SpamSieve - Re-Apply Filters** script. The effect of these two steps is that Mailsmith will download your mail quickly, and then filter it. This should be much faster than filtering the messages as they are downloaded.

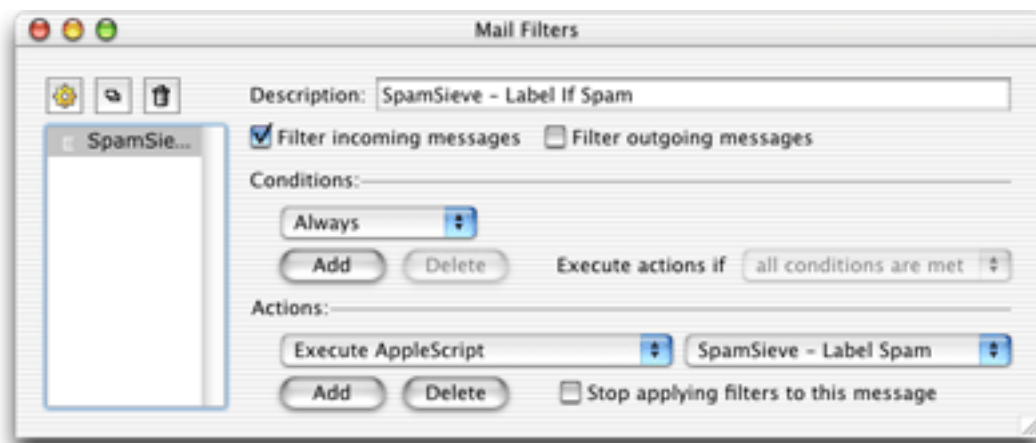
## 5.6 PowerMail

### 5.6.1 Basic Usage

- To add spam messages to SpamSieve, select one or more of them and run the **SpamSieve - Add Spam** script from the **Scripts** menu.
- To add good messages to SpamSieve, select one or more of them and run the **SpamSieve - Add Good** script from the **Scripts** menu.
- SpamSieve will label or move new spam messages that you receive, depending on how you configure PowerMail below.
- To manually label or move spam messages, select them and run the **SpamSieve - Label Spam** or the **SpamSieve - Move Spam** script from the **Scripts** menu.

### 5.6.2 Labelling Predicted Spam Messages

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



### 5.6.3 Moving Predicted Spam Messages

If, instead, you want SpamSieve to move suspected spam messages to a **Spam** folder, set up the filter to use the **SpamSieve - Move Spam** script instead.

## 5.7 Customizing the AppleScripts

### 5.7.1 Compatibility Notes

When editing an AppleScript, be sure to save it as a “compiled script” rather than a “Compiled Script.” The latter is a new format in Mac OS X 10.2 that some e-mail clients do not understand. The present beta of Script Editor 2.0 cannot save in the older, more compatible, format so you should use Script Editor 1.9 instead.

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

### 5.7.2 Moving Spam Messages After Adding Them

Normally when you add spam messages to the corpus, SpamSieve doesn’t move the messages in your e-mail client. If you prefer that the messages be moved to a **Spam** folder, you can do the following. Use Script Editor to open the **SpamSieve - Add Spam** AppleScript for your e-mail client. (Eudora users do not have such a script; they already have a command for adding and moving spam messages in the normal SpamSieve filter dialog.) The first line of the script contains the text:

```
property shouldMoveSpam : false
```

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

```
property shouldMoveSpam : true
```

Then save the script.

### 5.7.3 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](mailto:spamsieve@c-command.com)<sup>8</sup> so that I may assist you.

## 6 The Filter Menu

### 6.1 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:

---

<sup>8</sup><mailto:spamsieve@c-command.com>

**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).

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.

## 6.2 Remove Word...

You can remove words that you don't want in the corpus by selecting them and using the **Remove Word...** command in the **Filter** menu. Alternatively, the **Prune Corpus...** command will automatically remove many words that are not being used (see Section 6.3).

If you remove all the words from the corpus, for example using **Select All** followed by **Remove Word...**, SpamSieve will *also* forget all the messages that you added the corpus. This will enable you to retrain SpamSieve from scratch, and SpamSieve will let you use your old messages in the retraining.

## 6.3 Prune Corpus...

This command removes from the corpus words that are not contributing to spam recognition. This can dramatically decrease SpamSieve's memory use and launch time. You can prune repeatedly to reduce the corpus size further. However, excessive pruning will increase the training needed to recognize new types of messages.

## 6.4 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.

**Good Words** The number of unique words appearing in good messages.

**Spam Words** The number of unique words appearing in spam messages.

**Unused Words** The number of words that are not yet being used to identify spam messages, because they have not occurred enough times in your messages. When this number is very large, you may wish to prune your corpus (see Section 6.3) to reduce SpamSieve’s memory usage.

**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.

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.

## 6.5 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.

## 7 Preferences

### 7.1 Allow duplicate good/spam messages

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.

## 7.2 First check if message is in corpus

SpamSieve learns as you train it, but training is not instant. Adding a message to the corpus will not necessarily give SpamSieve 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 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.

## 7.3 Add predicted good/spam messages to corpus

These options cause SpamSieve to add messages to its corpus after it predicts whether they are good or spam. Thus, after the initial training you would only need to add messages to correct SpamSieve's mistakes. 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.

Adding predicted good messages tends to prevent false positives while slightly increasing false negatives. Adding predicted 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 add only predicted spam messages to its corpus. (Once you have achieved a good accuracy, you can even stop doing this, and from then on train SpamSieve only using its mistakes.)

## 7.4 Use system Address Book as a whitelist

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.

You can add addresses to the system address book using the Address Book application, or directly from an e-mail client that supports the system address book. You may want to add addresses to your whitelist without cluttering your address book with cards for each one. To do this, create a single person card in the address book using a name such as "Whitelist." Then add all the additional e-mail addresses to that one card.

This feature requires Mac OS X 10.2 or later.

## 7.5 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](mailto:spamsieve-fn@c-command.com)<sup>9</sup>. By looking at the messages that SpamSieve missed, I can improve its algorithms to catch such messages in the future.

## 8 Contact Information

You can download the latest version of SpamSieve from the [SpamSieve Web site](#)<sup>10</sup>. Questions about SpamSieve may be sent to [spamsieve@c-command.com](mailto:spamsieve@c-command.com)<sup>11</sup>. 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 Announcements mailing list](#)<sup>12</sup>. The traffic on this list is very low, only one message per new version of SpamSieve.

## 9 Registering

SpamSieve is shareware. If you find yourself using SpamSieve beyond a reasonable trial period, you must register it. Registration costs \$20 (US) and entitles you to free updates and support.

To register, choose **Registration...** from the **SpamSieve** menu. You can use the **Instant Registration...** button to enter your billing information from within SpamSieve or use the **Web Registration...** button to enter it at the [eSellerate Online Store](#)<sup>13</sup> in your Web browser.

Soon after paying, you'll receive an e-mail with your serial number. Enter it in the Registration 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 register again.

This is the honor system. If you use SpamSieve without registering, I probably won't know. However, registering will give me an incentive to continue updating and enhancing SpamSieve, and to write more Mac software. And you won't have to look at the shareware reminders anymore.

---

<sup>9</sup><mailto:spamsieve-fn@c-command.com>

<sup>10</sup><http://www.c-command.com/spamsieve/>

<sup>11</sup><mailto:spamsieve@c-command.com>

<sup>12</sup><http://www.c-command.com/spamsieve/support.shtml>

<sup>13</sup><http://store.eSellerate.net/mt/store>

## 10 Notes and Limitations

- Some users have reported that when SpamSieve is a login item it quits at launch with an error about an `NSPortTimeoutException`. To prevent this problem, simply do not make SpamSieve a login item. Your e-mail client will launch it automatically when it begins filtering messages.

## 11 Version History

### 1.3—February 11, 2003

- More resilient to spammers' trick 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 Emler 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 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.

# 12 Legal Stuff

## SpamSieve

SpamSieve and this manual are copyright © 2002–2003 by [Michael Tsai](#)<sup>14</sup>. All rights reserved.

Please distribute the unmodified `SpamSieve-1.3.dmg` file on the Web, LANs, compilation CD-ROMs, etc. Please do not charge for it (beyond a reasonable cost for media), or distribute the contents of the image file in isolation.

This software is provided by the copyright holders and contributors “as is” and any express or implied warranties, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall the regents or contributors be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

SpamSieve is a trademark of Michael Tsai. Mac is a registered trademark of Apple Computer. All other products mentioned are trademarks of their respective owners.

## EDCommon

[EDCommon](#)<sup>15</sup> 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

---

<sup>14</sup><mailto:mjt@c-command.com>

<sup>15</sup><http://www.mulle-kybernetik.com/software/EDFrameworks/download.html#EDCommon>

Doernenburg in all documents and publicity pertaining to direct or indirect use of this code or its derivatives.

THIS IS EXPERIMENTAL SOFTWARE AND IT IS KNOWN TO HAVE BUGS, SOME OF WHICH MAY HAVE SERIOUS CONSEQUENCES. THE COPYRIGHT HOLDER ALLOWS FREE USE OF THIS SOFTWARE IN ITS “AS IS” CONDITION. THE COPYRIGHT HOLDER DISCLAIMS ANY LIABILITY OF ANY KIND FOR ANY DAMAGES WHATSOEVER RESULTING DIRECTLY OR INDIRECTLY FROM THE USE OF THIS SOFTWARE OR OF ANY DERIVATIVE WORK.

## **EDMessage**

[EDMessage](#)<sup>16</sup> 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.

THIS IS EXPERIMENTAL SOFTWARE AND IT IS KNOWN TO HAVE BUGS, SOME OF WHICH MAY HAVE SERIOUS CONSEQUENCES. THE COPYRIGHT HOLDER ALLOWS FREE USE OF THIS SOFTWARE IN ITS “AS IS” CONDITION. THE COPYRIGHT HOLDER DISCLAIMS ANY LIABILITY OF ANY KIND FOR ANY DAMAGES WHATSOEVER RESULTING DIRECTLY OR INDIRECTLY FROM THE USE OF THIS SOFTWARE OR OF ANY DERIVATIVE WORK.

## **PCRE**

Regular expression support is provided by the [PCRE](#)<sup>17</sup> library package, which is open source software, written by Philip Hazel, and copyright by the University of Cambridge, England.

---

<sup>16</sup><http://www.mulle-kybernetik.com/software/EDFrameworks/download.html#EDMessage>

<sup>17</sup><http://www.pcre.org>