

The idea:

Professors want to write lecture notes for a course. We'll call that a book.

One chapter is the equivalent of a lecture. The simplest chapter will consist of text and code references into a single directory. The system will generate a single, static HTML page for viewing.

A book is marked by the paired tags `<BOOK></BOOK>`.

Chapters, sections, paragraphs, images, and code will be marked by paired tags:

<code><CHAPTER></CHAPTER></code>	<code><SECTION></SECTION></code>
<code><PARA></PARA></code>	<code><IMAGE></IMAGE></code>
<code><CODE></CODE></code>	

The features:

Each chapter has a title.

I load a chapter that contains several empty sections. Each section has a title.

I load a chapter with a section that contains several paragraphs of text.

I load a chapter that includes an image. The image is stored in a subdirectory of the directory that contains the chapter.

A chapter can contain a code example. The code is stored in a subdirectory of the directory that contains the chapter. The reference to the code includes the file name.

I load a chapter that contains a bullet list of several items.

I load a chapter in which images and code examples have captions as attributes.

Each of these "load" stories has a corresponding "generate a web page for" story.

A book can consist of one or more chapters.

A chapter can reside in a separate file rather than in text between its tags. In that case, the chapter tag has a FILENAME attribute:
<CHAPTER FILENAME="filename"></CHAPTER>.

I search a book for chapters that contain a certain string in their text. The program returns a collection of chapters.

I request a collection of all code examples in a book.

I generate an HTML page for a book that lists just its chapter titles.