

The art of ALT: toward a more accessible Web

John M. Slatin*

University of Texas at Austin, Austin, TX 78712, USA

Abstract

Continuing innovations in pedagogical uses of the Web are consistent with our discipline's long-standing commitment to the expansion of literacy. Surging interest in multimedia and visual rhetoric emphasizes the importance of the 1999 *Web Content Accessibility Guidelines* as a tool for instructors seeking to make their Web documents accessible to learners and colleagues who have disabilities. Text-only variants of media-rich sites are not sufficient; on the Web, as on our campuses, separate is not and cannot be equal. Changes in the way we approach designing class Web sites may be necessary to enable all learners to participate equally in the learning community. Accessibility is not a property of the document: It is situated in specific contexts and distributed across multiple agents and artifacts. A Web experience designed to be rich and meaningful for people with disabilities is likely to be rich and meaningful for those without disabilities as well; however, the reverse is not necessarily true. © 2001 Elsevier Science Inc. All rights reserved.

Keywords: Accessibility; ALT text; Americans with Disabilities Act (ADA); Disabilities; HTML; *Web Content Accessibility Guidelines*; Web design

1. Introduction

My purpose in this article is to persuade friends and colleagues in the computers and writing community to do more to ensure that the Web pages that they and their students compose—especially for classes—are accessible to people with disabilities. I raise this issue as a matter of some urgency in light of the surging interest in multimedia and visual rhetoric represented by this issue of *Computers and Composition*. I do not propose to swim against that current—we are clearly in the midst of a major shift in the ratio of image to text, as Richard Lanham (1993) and others have pointed out. But we must take care that in embracing

* Corresponding author.

E-mail address: jslatin@mail.utexas.edu (J.M. Slatin).

multimedia we do not disenfranchise students and colleagues with disabilities. This doesn't mean giving up an appeal to the senses; instead, it means learning to integrate a concern for accessible Web design into our practice as writers and teachers.

Several points crucial to this article are:

- that the purpose of a class Web site is not merely to present information but to enable participation in the learning community,
- that accessibility is not a property of the Web site, but is distributed among many interacting constituents,
- that an experience rich and meaningful for participants with disabilities is likely to be rich and meaningful for others also, and
- that text-only variants do not provide an equivalent alternative to media-rich sites that are not accessible to all students, faculty, and staff.

2. A civil rights issue

For almost 30 years, federal law has required that educational institutions make learning opportunities available to people with disabilities. This mandate has recently been extended to technology-enhanced teaching and learning. In 1998, the U.S. Department of Education's Office of Civil Rights held that colleges and universities must make electronically produced and disseminated materials as well as print-based materials available in accessible form (Waddell, 1999). A number of institutions—including the California Community College System with its 4.1 million students and the University of Texas at Austin with its 50,000 students—now have policies on the books committing them to provide accessible electronic materials.

3. The *Web Content Accessibility Guidelines*

The *Web Content Accessibility Guidelines* (hereafter WCAG; Chisholm, Venderheiden, & Jacobs, 1999) are the most important single resource for Web developers seeking to make their sites accessible. Published in May 1999 as a formal recommendation of the World Wide Web Consortium's Web Accessibility Initiative, the WCAG is the first of three separate but closely related documents that lay the groundwork for significantly enhancing the Web's accessibility. The WCAG deal with content as if it could be "device-independent"; the "Authoring Tool Accessibility Guidelines" (published as a formal recommendation in February 2000) describe how developers of Web authoring software should support efforts to produce accessible content and make their software accessible to authors with disabilities. Finally, the "User Agent Guidelines" (currently a working draft, most recently updated in October 2000) concern the software people use to access the Web, going well beyond familiar browsers to specialized tools such as speech-based browsers, mobile communications devices, and so on. The strategy of separating content from the software used to produce and access that content is theoretically vexed, but it is also rhetorically effective and technically pragmatic, enabling Web authors to understand how much we can do to improve accessibility even as we wait for authoring tools and user agents to catch up.

All three documents comprising the WCAG build on a foundation established in the Recommendation for HTML 4.0 (Raggett, Le Hors, & Jacobs, 1997). The principal changes from HTML 3.2 to 4.0 were introduced specifically to enhance accessibility (World Wide Web Consortium, 1999). This is also true for many of the changes in the specification for cascading style sheets, and for some of the features of SMIL, Synchronized Multimedia Integration Language (Jacobs & Brewer, 1999; Koivunen & Jacobs, 1999).

4. Accessibility is distributed

It should be clear by now that accessibility involves more than Web content alone. On the contrary, accessibility is *situated* in particular contexts and *distributed* across a number of interacting constituents. The experience of accessibility arises most immediately through the interaction of the document with the user agent (e.g., a Web browser). The quality of that interaction depends, in turn, on the extent to which authors heed HTML specifications and accessibility guidelines—and how individual users configure their browsers or other user agents. Finally, many people with disabilities use assistive and adaptive technologies of one sort or another, such as speech-based browsers (which read the page aloud) or voice-based ones (which accept voice commands). More common are screen-reading programs such as Freedom Scientific's JAWS and GW Micro's WINDOW-EYES, which read aloud material that appears on the screen. In the Microsoft WINDOWS environment, many of these technologies further depend upon the mediation of Microsoft's Active Accessibility, which passes information from applications like Microsoft INTERNET EXPLORER to the assistive device, which then presents it in a manner appropriate to the individual's needs.

I admit to a kind of perverse delight in describing this fiendish complexity. But making Web sites accessible doesn't necessarily mean keeping track of all these things from moment to moment. Indeed, the point is precisely that authors shouldn't have to. If authoring tools are designed in accordance with the authoring tool guidelines, for example, heeding the content guidelines will be far easier; if browsers comply with the user agent guidelines, people with disabilities will experience Web sites as richer and more meaningful.

5. Organization of the WCAG

In the absence of accessible authoring tools and user agents, 17 guidelines are included in the WCAG to help Web authors set priorities and plan for accessibility. Under each guideline is a set of checkpoints, arranged according to three priority levels. **Priority 1** items are essential to accessibility; failure to address priority 1 issues guarantees that some people will be unable to access some features of the site. **Priority 2** refers to items that are important to accessibility; failure to attend to priority 2 concerns means that some people will have serious difficulty with the site. **Priority 3** refers to choices that enhance accessibility for some users. These are the kind of little touches that may not even be noticeable individually, but make the experience a more pleasant one.

Included in the WCAG is also a checklist organized according to the HTML elements that

also might appear on a given site. If a site uses images, for example, the designer can refer to the image checkpoints to see what practices fall under priority 1, priority 2, or priority 3 and then follow the ones that apply. By contrast, if the site contains no images, the designer need not address that set of checkpoints. A “techniques” document in the WCAG provides detailed examples of how to implement each checkpoint.

6. The Texas 2000 Living Museum as an example

The remainder of this article describes the work of me and my colleagues at the Institute for Technology and Learning as we designed the 2000–2001 “edition” of an ongoing Web-based project, TX2K: The Texas 2000 Living Museum (Slatin et al., 2000). TX2K is a year-long, interdisciplinary project designed for teachers and students throughout Texas. Students act as curators of the living museum, researching and preparing three exhibits about their communities and publishing them on the Web site. Carefully structured research activities combine local materials with digital resources in the University of Texas collections and elsewhere on the Web. Students write day-in-the-life narratives to dramatize their exhibits. Virtual field trips and online dialogues link participants with peers in other communities. Summer workshops prepare teachers for successful participation.

With over 30 schools set to participate in 2000 and 2001, TX2K is a large, complex, and highly interactive site. Besides two sets of online documentation, there are more than 30 different page layouts in six major components: Lobby, Gallery, Curator’s Workshop, Teacher’s Office, Library, Message Boards. There are four different views of the site (Guest, Student, Teacher, ITAL staff) corresponding to different levels of access.

Our key design goals were to make the site significantly easier to learn and use, significantly faster, and more fully graphical than it had been in previous years, while more fully honoring our commitment to guarantee access for learners, parents, and teachers with disabilities. At first these goals seemed incompatible. What rescued us, I think, was deciding that meeting the needs of learners with disabilities was our starting-point for design. From this it followed that we should treat accessibility requirements not merely as constraints, but as functional requirements to be exploited esthetically, features treated like design elements, such as wheelchair ramps in buildings after the Americans with Disabilities Act (ADA) or the conduits and stairways in the Beaubourg Museum in Paris. This approach, which I call *access-first design*, enabled the seemingly incompatible goals outlined a moment ago to cooperate and support one another.

7. Access-first design

The access-first concept is a logical and pragmatic extension of user-centered approaches to software development, and also of a personal and disciplinary commitment to student-centered pedagogies. As teachers and Web developers, we aim to create esthetically rich, intellectually exciting learning experiences that engage and challenge all students to the full extent of their abilities. The goal of accessible design is to enable people who have disabilities to participate as equals in those learning experiences. An experience rich and

meaningful for students (or teachers) with disabilities is likely to be rich and meaningful for those without disabilities, though the reverse may not be true. Access-first design, then, offers a way to enhance the experience for everyone.

7.1. Enabling participation

To support participation in and identification with the learning community, we must enable participants to recognize, understand, and remember how the site is organized. It is critical, therefore, to:

- devise a clear, consistent navigation scheme (WCAG #13),
- provide contextual and orientation information (WCAG #12), and
- provide an “equivalent alternative” for any visual or auditory material (WCAG #1).

Meeting these requirements has the advantage of making the site easier to learn and use, and faster to navigate as well.

7.2. Clear, consistent navigation

Site architecture is critical to accessibility. Organizing TX2K as a suite of virtual rooms, each dedicated to specific functions, serves as an aid to memory for participants as. The Curator’s Workshop, for example, provides facilities for creating and working on exhibits, and the Museum Gallery, shown in Figure 1 below, displays student work. Internal links point to tools available in the current room, while external links lead to other rooms in the museum. Graphics also play an important role in helping visual learners understand how to navigate the museum. Each room is clearly labeled; external links are placed along the “rear wall” of each room; internal links are closer to the foreground.

7.3. Context and orientation: page titles and link text

Writing page titles that identify unique pages and indicate their relation to other parts of the site (TX2K Library: Infrastructure Resources, for example) helps participants refine their mental maps of the site. Matching page titles with link text, both on the page and in ALT attributes, is also helpful. For example, a hyperlink that reads “Go to the Library” should lead to a page called “TX2K Library” instead of a page titled “Find a resource.” Most participants would find it fairly easy to connect the idea of going to the library with finding a resource, but people with cognitive impairments might find this unnecessarily difficult, and the same might be true for participants for whom the site is an entirely auditory experience.

8. The art of ALT

The “prime directive” of the WCAG is to “provide equivalent alternatives for all visual and auditory material” (online). The principal tool for this purpose is commonly referred to as the ALT tag or ALT text. The ALT text is a phrase or sentence “attached” to an image or other element so that people who use assistive technologies such as screenreaders and talking browsers—or people who, for the sake of speed, prefer to browse the Web with



Fig. 1. The TX2K Gallery (Teacher's View). External links lead to the Curator's Workshop and the Teacher's Office. Internal links point to search tools and other information.

images turned off—can identify the element. For most people, ALT text typically appears briefly when the mouse-pointer passes over an image or other element that is tagged in this way; these users will see a small box containing a sentence or phrase that disappears after a second or two.¹

ALT text should do two things: (1) briefly identify the nontextual element to which it is attached, and (2) provide access to the *functionality* represented by that element. The ALT attribute is supposed to be equivalent to the element it's attached to, but participants who need ALT text are usually in no position to judge equivalence because they don't encounter the original in the first place. ALT text must therefore be succinct, descriptive, and accurate. The WCAG do not specify how long ALT text should be, and indeed the theoretical upper limit is over 65,000 characters (Maden, 1998). But assistive technologies may impose *de facto* limits. An obscure setting for the JAWS screenreader causes problems if ALT text exceeds 150 characters, for instance. Although individual users can change this setting, few know how, so it is probably wiser to accept a 150-character limit as a constraint and a spur to creativity.

Writing effective ALT text, then, is an exercise in extreme economy. The challenge is to include as much information as possible in the fewest characters without sacrificing intelligibility. In designing ALT text for the TX2K Museum Gallery, for example, we concluded that “search for exhibits by city, school, or exhibit name by selecting the search globe” (82 characters) would be more useful than the much stingier “search” because it alerts partici-

pants in advance that different ways of searching are available. Note that participants who don't want to listen to the whole ALT text can activate the link or move forward as soon as they're ready, similar to the way a good voice-mail system (or is that a contradiction in terms?) supports type-ahead.

9. Listening to the Web

ALT text and other contextual cues that are easy to understand and remember are critical tools for constructing accurate cognitive maps for users who experience a site aurally. ALT tags have to make sense when read aloud, both one at a time and in a sequence. It's very rare that they do; often the tags exist in a jumble.

To achieve auditory coherence, it may be necessary to rethink the relationship between text and graphics, and even to reorient the design process. HTML is structured in such a way that ALT text is clearly subordinate to the image. It is by definition an attribute of the IMG (image) or other HTML element, and cannot stand on its own. This encourages the practice of designing graphics and laying out the page for visual appeal first, leaving the composition of ALT text and long descriptions to the very end. The almost inevitable result is a page that sounds like an incoherent laundry list of disconnected items when read aloud.²

These problems arise even when developers are sensitive to the issue. TX2K's major pages, for example, are organized for auditory presentation in such a way that the first thing the participant hears is the name of the room, followed by a link to a page containing an extended prose description of the room and its features. This "d-link" is followed by a link to the navigation bar at the bottom of the screen, enabling participants to jump quickly to another part of the site if they find themselves in the wrong place. Tags that identify the room and point to the navigation bar are associated with transparent GIFs; the link to the detailed description of the room is associated with a stylized "d" in the upper left corner of the screen.

So far, so good. But the next two links in the auditory sequence refer to pages other than the Museum Gallery—the Curator's Workshop and the Teacher's Office. This design is not a problem visually: the links are clearly contextualized by the graphic (see Figure 1). But for someone listening to the page, being directed to other parts of the site rather than hearing about what is available "locally" may be disconcerting.

This sort of problem is an artifact of the conventional, graphics-first design process we followed in developing TX2K. One remedy is to treat the page as a coherent *text*. We write out the ALT tags and read them aloud, revising both the individual phrases and the entire sequence of tags until each individual tag and the page as a whole make sense. Then we lay out the page visually, arranging individual graphics to force the ALT tags into the appropriate reading order. This approach has a number of advantages, not least of which is that it reasserts the importance of written language in the context of multimedia composition.

10. Separate is not equal: against text-only variants

You may be asking if it wouldn't be easier just to put up a text-only variant of the site, as so many people seem to be doing now. Yes, it would be easier—but, as Richard Nixon

used to say, it would be wrong. It's not just that maintaining two separate versions of a large, complex site is expensive and difficult. As the Supreme Court held nearly 50 years ago, separate is not and cannot be equal. It is virtually impossible *not* to think of the media-rich variant as the "real" and therefore privileged site—the "original" to the text-only ALT version—and to pay more attention to the "real thing." The dual versions will fall out of synch, like Braille menus at hotel restaurants. How many users would be likely to replace the media-rich site with the text-only variant?

11. Conclusion: boundary objects

The first evidence that the access-first approach may work came in the form of generous support from Microsoft and an Extraordinary Web Design award from Project EASI³ acknowledging our combination of a richly graphical interface with full accessibility. But TX2K is only a beginning.

Web documents encode the social practices from which they emerge and to which they give rise (Brown & Duguid, 1995, 2000). All of these materials—the guidelines, the checklist, the techniques documents, this article—are boundary objects, reifying the results of extended discussion and debate among different communities of practice. *Reification* tends to be a dirty word in English Studies. But reification and participation form a duality, a complementary pair (Wenger, 1998) whose interplay is an inevitable and essential part of any serious, collective effort to express complex ethical, legal, moral, political, and social values in technical terms (Friedman, 1998). As rhetoricians of the Web and teachers of writing in webbed environments we can now participate in reshaping practice.

Notes

1. Technically speaking, ALT text is defined as an *attribute* of an HTML *element*. (Elements can stand alone; attributes, as the name suggests, cannot.) For example, HTML defines an image as an element called IMG; ALT text is treated as an attribute of the IMG element. Let's suppose that I want to put a picture of my head on my Web page. The picture is stored in a file called "myhead.jpg." Without ALT text, people using talking browsers or screenreaders would hear only the actual filename. But if I supply an ALT attribute, users will hear the ALT text instead. The syntax looks like this in the HTML document: ``. Note that this image has several attributes, including height, width, and a "source" file, as well as the ALT text.
2. A *long description* (LONGDESC) attribute was introduced in HTML 4.0 to enable Web authors to provide extended descriptions of complex images, graphs, charts, and so on, but this attribute is not yet supported by the major commercial browsers.
3. Project EASI (Equal Access to Software and Information) was established in the late 1980s by Professor Norman Coombs, an historian at the Rochester Institute of Tech-

nology, after he lost his sight in an automobile accident. EASI maintains a comprehensive Web site at <<http://www.rit.edu/~easi>> addressing many aspects of access to information technology.

Acknowledgment

ITAL's efforts to enhance TX2K's accessibility have been supported by a generous grant from the Microsoft Corporation. TX2K is a team effort: thanks to Tiffany Elliott, Elizabeth Flammang, Josh Knowles, Elizabeth Lehnertz, Cristy Oliver, Jay Overfield, Chet Salazar, Annie So, Ross Speir, and Vennila Thurimavalavan.

John Slatin is Director of the Institute for Technology and Learning at the University of Texas at Austin, and former director of the University's Computer Writing and Research Lab.

References

- Brown, John Seely, & Duguid, Paul. (1995). The social life of documents. Available: <http://www.parc.xerox.com/ops/members/brown/papers/sociallife.html> [Accessed: November 9, 2000].
- Brown, John Seely, & Duguid, Paul. (2000). *The social life of information*. Cambridge, MA: Harvard Business Press.
- Chisholm, Wendy, Vanderheiden, Gregory, & Jacobs, Ian. (Eds.). (1999). Web content accessibility guidelines 1.0. World Wide Web Consortium Web Accessibility Initiative. Available: <http://www.w3.org/TR/WAI-WEBCONTENT/> [Accessed: July 15, 2000].
- Friedman, Batya. (Ed.). (1998). *Human values and the design of computer technology*. CSLI Lecture Notes, No. 72. New York: Oxford University Press.
- Jacobs, Ian, & Brewer, Judy. (Eds.). (1999). Accessibility features of CSS. World Wide Web Consortium. Available: <http://www.w3.org/TR/CSS-access> [Accessed: July 15, 2000].
- Koivunen, M. R., & Jacobs, Ian. (Eds.). (1999). Accessibility features of SMIL. World Wide Web Consortium. Available: <http://www.w3.org/TR/SMIL-access/#Introduction> [Accessed: July 15, 2000].
- Lanham, Richard A. (1993). *The electronic word: Democracy, technology, and the arts*. Chicago and London: Chicago University Press.
- Maden, Chris. (1998). Re: D-link and LONGDESC (GL type stuff). April 23, 1998. WAI-IG mailing list. World Wide Web Consortium. Available: <http://lists.w3.org/Archives/Public/w3c-wai-ig/1998AprJun/0058.html> [Accessed: July 16, 2000].
- Raggett, David, Le Hors, Arnaud, & Jacobs, Ian. (1997). HTML 4.0 specification. World Wide Web Consortium. Available: <http://www.w3.org/TR/REC-html40/> [Accessed: November 9, 2000].
- Slatin, John, M., et al. (2000). TX2K: the Texas 2000 living museum. Austin, Texas: Institute for Technology and Learning. Available: <http://www.ital.utexas.edu/tx2k> [Accessed: July 15, 2000].
- Waddell, Cynthia, D. (1999). The growing digital divide in access for people with disabilities: Overcoming barriers to participation. Available: <http://www.aasa.dshs.wa.gov/access/waddell.htm> [Accessed: July 18, 2000].
- Wenger, Etienne. (1998). *Communities of practice: learning, meaning, and identity. Learning and doing: Cognitive, social, and computational perspectives*. New York: Cambridge University Press.
- World Wide Web Consortium. (1999). HTML 4.0 accessibility improvements. Available: <http://www.w3.org/WAI/References/HTML4-access> [Accessed: July 15, 2000].