5
Experiments

In this chapter, we describe an extensive evaluation of our spoken content publishing technique using real Web video files from the educational domain. Besides evaluating the efficacy of the publishing technique, our goals included an in-depth analysis of the hits on the Web pages synthesized for the video files.

5.1
Experimental Setup

We evaluated the tool using a set of 99 learning objects (LO's) consisting of 10-minute long video files, with dialogs covering elementary Chemistry topics.

We conducted the experiments in two stages, for 8 months. The first stage lasted 5 months and all 99 LO's were published using content-independent metadata, i.e., "*Title*", "*Author*", "*Keywords*" and "*Short description*". The goal of the first stage was to equalize the Web page hits using content-independent metadata and use this information to create two balanced groups of LO's to assess the efficacy of the tool. We sorted LO's by the number of hits and then, for each pair of LO's whose number of hits have the same order of magnitude, we assigned each of them to a different group.

The second stage lasted 3 months. We submitted one of the groups to the tool, but not the other. The goal of this second stage was to evaluate the efficacy of the tool. We hosted the static Web pages on the WordPress server and the video files on Youtube and used the statistics these services provide to evaluate the tool.

In what follows, let Group *A* refer to the set of LO's published in the first stage, Group P refer to the set of LO's published using the tool in the second stage, and Group ¬P refer to the set of LO's described only by content-independent metadata.

5.2
Data Analysis

We present in this section the results of the hit analysis for both stages of the experiment.

5.2.1
An analysis of total number of hits

Approximately 75K hits were obtained during both stages of the experiment. However, Group A obtained just 22% from the total number of hits.

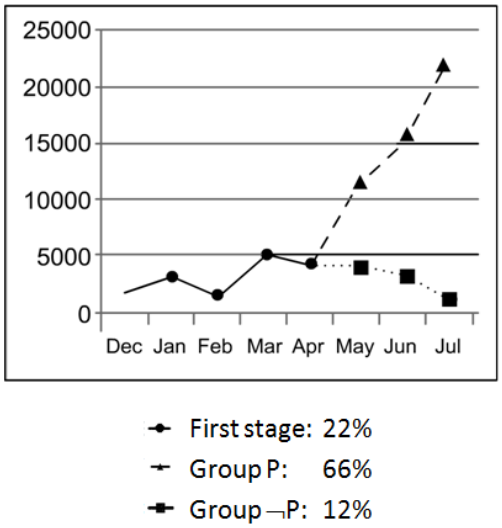

Figure 27 Hit stats. First stage 1–5. Second stage 5–8.

Note that Group A represents the first stage with all 99 LO's. Also note that first stage collected data during 5 months, whereas the second stage lasted 3 months. Hence, 78% of the hits were performed in the second stage (see Figure 27). Observe that Group P obtained 66% of the total number of hits and Group ¬P just 12%. Thus, Group P obtained 84% of the total number of hits in the second stage, i.e., 5.3 times more than the number of hits obtained by Group ¬P.

5.2.2
An analysis of the number of hits by language

As described in Section 3.3, the translation step created a new static Web page, for each new language. In the experiment, three scripts were generated for each asset, in English, Spanish and Portuguese. Figure 28 gives, for each

translated static Web page, the percentage of total number of hits: 43% for pages in Spanish, 37% for pages in Portuguese and 20% for pages in English. Note that 63% of the hits came from the translated Web pages.

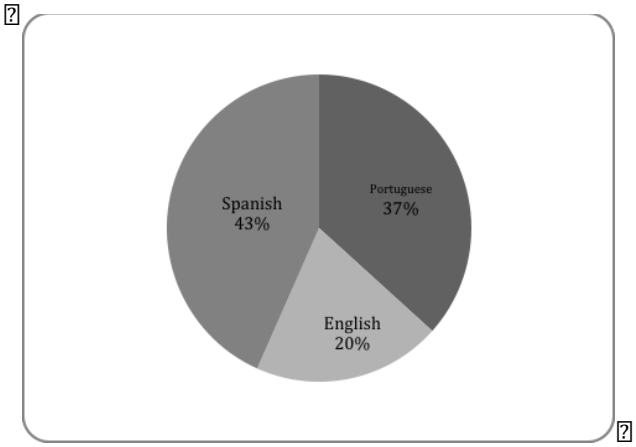

Figure 28 Hits percentage of translated static Web pages generated by our publishing technique. 37% of the hits were obtained by the Web pages of the assets in their native language and 67% of the hits were obtained by the translated Web pages.

Although the visitors of our static Web pages are anonymous, i.e., they do not provide information about where they come from or his/her mother language, YouTube provides information about logged in users. The information obtained from YouTube is very important because all actions (share, comment or mark as favorite) are done by users that are somehow interested in the content of an asset.

Using such information, Figure 29 indicates that only users from Brazil, Portugal and Japan shared, commented or marked as favorites the assets of Group A.

After applying the technique, creating Group P, users from other countries (Sweden, Japan, Brazil, Spain, Peru, United States and Portugal) could be reached, as shown in Figure 30. Note that the native language of Brazil, Spain, Peru, United States and Portugal is indeed English, Spanish or Portuguese, the languages available through our technique. Curiously enough, this is not true of Japan and Sweden, although there is a sizable population of Brazilians living in

Japan and English is the second most spoken language in both of these countries.

Figure 29 Countries that have interacted with the content: Japan, Brazil and Portugal.

Figure 30 Countries that have interacted with the content: Sweden, Japan, Brazil, Spain, Peru, USA and Portugal.

LO's with the largest number of hits	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Group ¬P	1636	1615	1416	969	812	642	514	300	265	220
Group P (using the tool)	2466	1774	1744	1562	1499	1467	1421	1414	1386	1307

Table 2 Top 10 LO's by number of hits.

Least hit LO's	10	9	8	7	6	5	4	3	2	1
Group ¬P	32	27	27	25	24	21	21	12	11	7
Group P (using the tool)	333	326	287	271	250	213	195	191	108	95

Table 3 Last 10 least hit Learning Objects.

5.2.3
An analysis of the number of hits by asset

This part of the experiment addressed the question: What is the difference between the number of hits of an LO that was published using the tool and an LO described by content-independent metadata? Table **2** shows the top 10 LO's, according to the number of hits, where the second line corresponds to LO's in Group ¬P and the third line to LO's in Group P. Table **2** indicates that the number of hits for LO's indexed by the tool (Group P) is indeed greater than the number of hits for LO's described only by content-independent metadata (Group ¬P), as expected. This is true in general for all 99 LO's. Table 3 shows the LO's with the least number of hits, which have an even greater discrepancy.

5.2.4
Regional analysis

Since Youtube provides the user's location, it was possible to tabulate the number of hits by continent (with the exception of Oceania).

In the first stage of the analysis, Group A obtained 13,911 hits from South America, 73 hits from North America, 39 hits from Asia, 1,297 hits from Europe and 31 hits from Africa. We note that, from the hits in South America, 13,673 (approx. 99%) were from Brazil, which is a Portuguese-speaking country. The same was observed in Europe where, out of the 1,297 hits, 1,107 (approx. 85%) were from Portugal, which is also a Portuguese-speaking country.

In the second stage, the assets obtained 52,366 hits from South America, 779 hits from North America, 494 hits from Asia, 3,738 hits from Europe and 167 hits from Africa. We again highlight that, out of the number of hits from South

America, most were from Brazil. However, during a short period, the number of hits from the rest of South America, mostly from Spanish-speaking countries, increased from 1% to 5.5%. (Note that the population of Brazil is almost 50% of the population of South America). The number of hits from Europe was also less concentrated: Portugal, that had 85% of the hits in the first stage of the experiment, had 67% in the second stage; the total number of hits from other European countries more than doubled from 15% to 33%.

Figure 31 shows the ratio increase by continent, obtained by dividing the number of hits in the second stage of the experiment by the number of hits in the first stage.

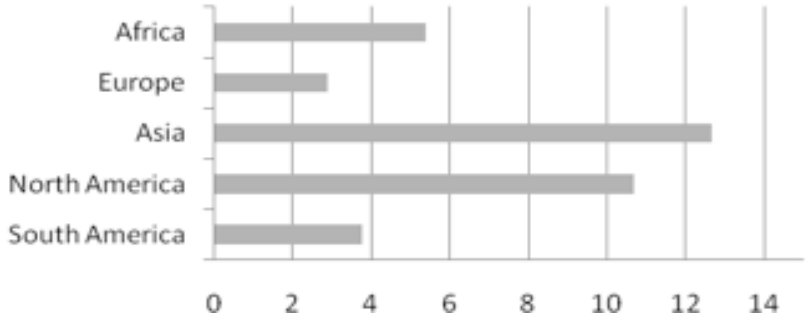

Figure 31 The number of hits boosted in different orders of magnitude for each continent.