

Figure 5. Level of difficulty for loop creation rated by the subjects on a scale of 0 to 5 (5 being extremely difficult). The labels are defined as — S1 and S2 are stage 1 and 2 of manual loop creation, LM-A and LM-S are LoopMaker using automatic and semi-automatic mode.

than manual loop creation. We then performed independent listening tests to compare the quality of the loops created using our system to the manually created loops.

User Study

For this study, we recruited 14 participants from the age of 24 to 53. The participants were a mixture of experts (6) and novices (8). We define an expert as someone who claimed to be both knowledgeable in the use of audio editing software and possess at least a basic music theory background (i.e. play a musical instrument or having taken a music theory class). We define all non-expert participants as novices. We asked each participant to manually create three loops from each of three different songs (total of nine loops). They were allowed to use any audio editing software package they like for this creation. We asked participants with no prior audio editing experience to use Audacity. They were allowed to use any online tutorial of their choice in order to gain a basic understanding of Audacity and audio editing. Two participants were not able to complete the task, as they found the task and use of the software too challenging, but all other participants completed the task.

We asked each participant to rate the difficulty of creating loops on a scale of 0 to 5 (0 referring to extremely easy and 5 referring to extremely difficult). We also asked them to individually rate the two stages of the process. The first stage is to listen to the song and roughly find the location of each loop (steps 1 and 2 mentioned in Section 1). The second stage is to refine the boundaries and create the loop (steps 3 and 4 mentioned in Section 1). The first stage also corresponds to the user input when using the semi-automatic mode of operation.

We then asked each participant to automatically create loops using LoopMaker with both modes and rate the level of difficulty (using the above scale). The music files we used in our experiment were instrumental music files that were randomly selected from the RWC Music Database [12].

We noted the time taken to complete each task, the reported difficulty levels, and comments left by the participants. The length of the manually created loops varies from 3 seconds to 13 seconds with an average of 5.2 seconds. Both novices

and experts take about 1 minute per loop to find the approximate location (stage 1). To refine the location of the loops (stage 2), novices took a median of 7.5 minutes and experts took a median of 3 minutes per loop. The time taken for this task using LoopMaker in both automatic and semi-automatic (after completing stage 1 manually) modes is within milliseconds as it happens at the press of a button and has very low computational complexity.

As shown in Figure 5, the median level of difficulty reported by novices was 4 for manual loop creation (3 for stage 1 and 4.25 for stage 2). The median level of difficulty reported by experts was 3 (2 for stage 1 and 4 for stage 2). This is consistent with the amount of time taken for the same things. The difficulty in creating loops using LoopMaker (both automatic and semi-automatic modes) was reported to be significantly lower.

Some novices preferred using our system to find all of the loops automatically, and most of the participants felt that our semi-automatic loop creation is very helpful because it automates stage 2, which both novices and experts found to be difficult. Some experts found it helpful to see the beat location in their favorite music editing software package during stage 2. However, they reported that the task was still difficult.

We found that several experts use the following general procedure to perform the task — they preserve regularity of beats, then detect a transition point by comparing similarity of loudness, consistency in instruments playing and similarity in harmony. This corresponds closely to the way our algorithm functions.

Listening Tests

We compared the quality of the manually created loops to those that were created by LoopMaker, using listening tests. In order to perform an unbiased comparison, we had third parties perform the comparison using Amazon Mechanical Turk (AMT). We define the quality of a loop in terms of how natural and seamless it sounds. This is to say that a loop that sounds like continuous music without unnaturally abrupt transitions when played repeatedly is considered to be a high quality loop. Given the inherently ambiguous nature of this definition, we performed three different kinds of listening tests to judge quality.

The total number of manually created loops was 108 (12 participants that completed the task \times 9 loops/participant). LoopMaker can automatically create a large number of loops per song. However, in order to make the comparison as comparable as possible, we use automatically created loops that were at about the same location as the manually created loops. In order to do this, for each manually created loop, we automatically found the corresponding location in the given song using dynamic time warping [16], and then used the semi-automatic mode of LoopMaker to find the closest possible loop. We therefore have a loop that was created by LoopMaker corresponding to each of the 108 manually created loops.

We recruited 1337 subjects on AMT. We screened each subject with a listening test as was previously done for other AMT based audio evaluation tasks [7]. Specifically, they were asked to listen to three audio segments. Each segment had a different

Type	Novices	Experts	Overall
LoopMaker	3.864	3.759	3.786
Manual	3.358	3.788	3.517

Table 1. Average naturalness rating on a scale of 1(Very Abrupt) to 5(Very Natural).

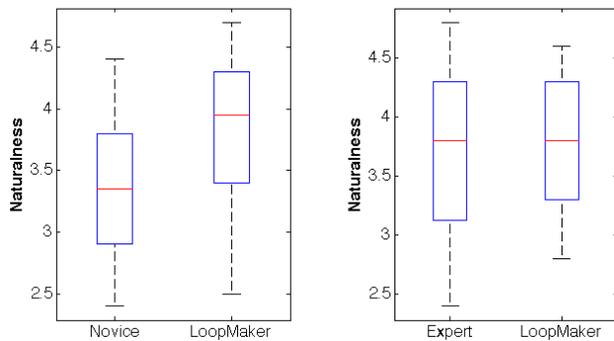


Figure 6. Naturalness rating.

number of sine tones of varying frequencies. They were asked to report the number of sine tones. If the reported number of tones was off by more than one, we discarded the response of the given subject. All but four subjects passed the screening.

We asked each of these subjects to compare 6 manually created loops with the 6 corresponding automatically created loops. They were randomly provided with 6 out of the 108 loops. Below, we describe the three listening tests that were performed by each subject. For each test, we presented the subject with the 12 loops mentioned above. Moreover, we played each loop continuously four times so that they can hear each transition multiple times. This task took each subject a total of about 10 minutes and we paid them for \$1.50 for the task.

Naturalness Rating

The goal of this test was to determine how natural the transitions sound to subjects. We asked each subject to rate how natural each loop sounds on a scale of 1 to 5 where 5 means very natural, 4 means moderately natural, 3 means acceptable, 2 means not so pleasant, and 1 means very abrupt. We specifically mentioned that changes in sections or parts of music do not qualify as abrupt changes.

As shown in Table 1, the loops that were created by LoopMaker were rated higher than those created by novices (t-statistic:3.3526, p-value:0.001415) and comparable to those created by experts (t-statistic:2.4579,p-value:0.01525). However, as shown in Figure 6, the variance for those created by LoopMaker is less than the variance of those created by experts.

Abruptness Rating

The goal of this test was to determine if subjects could hear abrupt changes. For each of the 12 loops, we asked them click a button when they hear a change. They were allowed to click the button multiple times per loop. We simply tallied

Type	Novices	Experts
LoopMaker	41.19%	45.25%
Manual	57.07 %	57.97%

Table 2. Percentage of abrupt changes detected.

Natural Preference	Novices	Experts	Overall
LoopMaker	41%	39%	41%
Manual	31%	30%	29%
Both Natural	16%	17%	17%
Both Abrupt	12%	14%	13%

Table 3. Preference of naturalness.

the total number of abrupt changes that they could hear and report this as percentage of the total number of transitions. We hypothesize that they would hear fewer abrupt changes in a more natural sounding loop.

As shown in Table 2, the subjects reported less abruptness in loops that were created by LoopMaker when compared to those created by both novices (t-statistic:-2.3843, p-value:0.02042) and experts(t-statistic:-2.1512, p-value:0.03465).

Comparison to Manual Loops

The goal of this test was to directly compare loops that were created by LoopMaker to manually created loops. We presented each subject with each of the 6 pairs of corresponding loops. For each pair, we asked the subject to select one of the following: Clip A sounds more natural, Clip B sounds more natural, Both sound natural, and both sound abrupt.

As shown in Table 3, on average, subjects preferred loops that were created by LoopMaker to those created by both novices and experts.

All three listening tests suggest that loops that were created by LoopMaker are of higher quality than those created by novices and at least comparable to those created by experts. Note that we used LoopMaker to create loops that were comparable in location to those that were manually created. We did not use the loops with the highest confidence score. We argue that a real-world use case would not have this inherent bias against our method, and believe that users would be more likely to use loops with higher confidence scores. Therefore, we think that the results would be of even higher quality in practice. To test this, we examined the subset of loops generated by LoopMaker that received a naturalness rating of 4 or 5 and found that only 25% of the transitions in these loops had abrupt transitions, whereas 57% of the transitions from manually created loops were abrupt.

Finally, we simply asked each participant in the study if they feel that our system is useful and if they would like to use it. 80% of the participants answered this question positively.

CONCLUSION

We have presented an interactive system to assist novices and experts in creating music loops. Our user studies helped

identify tedious and challenging parts of the loop creation process. Our system automates these parts, allowing users to concentrate on the more creative aspects of the process. These studies suggest that loops created using our system are significantly easier to create and are of comparable quality to those created by experts. We aim to explore other aspects of audio content creation that enable users to concentrate on the creative aspects of the process.

ACKNOWLEDGMENTS

We thank our colleagues at Adobe Research for participating in our user study, and Karrie Karahalios for providing valuable feedback on our paper.

REFERENCES

1. Jim Aikin. 2001. Making loop music in a computer. *Keyboard* 27, 2 (02 2001), 134–135.
2. Luca Benedetti, Holger Winnemöller, Massimiliano Corsini, and Roberto Scopigno. 2014. Painting with bob: Assisted creativity for novices. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. ACM, 419–428.
3. William E Benjamin. 1984. A theory of musical meter. *Music Perception: An Interdisciplinary Journal* 1, 4 (1984), 355–413.
4. Adam Berenzweig, Beth Logan, Daniel PW Ellis, and Brian Whitman. 2004. A large-scale evaluation of acoustic and subjective music-similarity measures. *Computer Music Journal* 28, 2 (2004), 63–76.
5. Nicholas J. Bryan, Gautham J. Mysore, and Ge Wang. 2014. ISSE: an interactive source separation editor. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM.
6. J. Brzezinski. 2009. Game sound: an introduction to the history, theory, and practice of video game music and sound design. *Choice* 46, 7 (03 2009), 1323.
7. Mark Cartwright, Bryan Pardo, Gautham J Mysore, and Matt Hoffman. 2016. Fast and easy crowdsourced perceptual audio evaluation. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*. IEEE, 619–623.
8. Matthew E.P. Davies, Philippe Hamel, Kazuyoshi Yoshii, and Masataka Goto. 2014. AutoMashUpper: Automatic Creation of Multi-Song Music Mashups. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 22, 12 (2014), 1726–1737.
9. Nicholas Davis, Holger Winnemöller, Mira Dontcheva, and Ellen Yi-Luen Do. 2013. Toward a cognitive theory of creativity support. In *Proceedings of the 9th ACM Conference on Creativity & Cognition*. ACM, 13–22.
10. Daniel PW Ellis. 2007. Beat tracking by dynamic programming. In *Journal of New Music Research*, Vol. 36. 51–60.
11. Bill Gibson. 2005. *The S.M.A.R.T. Guide to Producing Music with Samples, Loops, and MIDI*. Artistpro, Boston, MA, USA.
12. Masataka Goto, Hiroki Hashiguchi, Takuichi Nishimura, and Ryuichi Oka. 2002. RWC Music Database: Popular, Classical and Jazz Music Databases.. In *ISMIR*, Vol. 2. 287–288.
13. Tatsunori Hirai, Hironori Doi, and Shigeo Morishima. 2015. MusicMixer: computer-aided DJ system based on an automatic song mixing. In *Proceedings of the 12th International Conference on Advances in Computer Entertainment Technology*. ACM, 41.
14. Ning Hu, Roger B Dannenberg, and George Tzanetakis. 2003. Polyphonic audio matching and alignment for music retrieval. In *Applications of Signal Processing to Audio and Acoustics, 2003 IEEE Workshop on*. IEEE, 185–188.
15. Tetsuro Kitahara, Kosuke Iijima, Misaki Okada, Yuji Yamashita, and Ayaka Tsuruoka. 2015. A loop sequencer that selects music loops based on the degree of excitement. *Proceedings of the 12th International Conference in Sound and Music Computing, SMC 2015 (2015)*, 435–438.
16. Meinard Müller. 2007. Dynamic time warping. In *Information Retrieval for Music and Motion*. 69–84.
17. Bee Suan Ong and Sebastian Streich. 2008. Music loop extraction from digital audio signals. In *IEEE International Conference on Multimedia and Expo*. 453–469.
18. David E. Reese, Lynne S. Gross, and Brian Gross. 2009. *Audio production worktext: concepts, techniques, and equipment*. Taylor & Francis.
19. Steve Rubin, Floraine Berthouzoz, Gautham J. Mysore, Wilmot Li, and Maneesh Agrawala. 2000. Content-based tools for editing audio stories. In *Proceedings of International Society of Music Information Retrieval*. ACM, 113–122.
20. Joan Serra, Emilia Gómez, Perfecto Herrera, and Xavier Serra. 2008. Chroma binary similarity and local alignment applied to cover song identification. *IEEE Transactions on Audio, Speech, and Language Processing* 16, 6 (2008), 1138–1151.
21. Sebastian Streich and Bee Suan Ong. 2008. A Music Loop Explorer System. In *International Computer Music Conference Proceedings*.
22. Adobe Creative Team. 2012. Extending and Shortening Musical Selections in Adobe Audition CS6. (2012). <http://www.adobepress.com/articles/article.asp?p=1867758&seqNum=7>.
23. Newspapers Tribune. 2004. Making music; Software makes it easy for amaterus to compose tunes. (04 2004).
24. Brian Whitman. *The Infinite Jukebox*. <http://labs.echonest.com/Uploader/index.html>.
25. Jun Xie, Aaron Hertzmann, Wilmot Li, and Holger Winnemöller. 2014. PortraitSketch: Face sketching assistance for novices. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. ACM, 407–417.