

Region Radio: An AI that Finds and Tells Stories about Places

Douglas H. Fisher Emily Markert Abigail Roberts Kamala Varma

Department of Electrical Engineering & Computer Science

Vanderbilt University

Nashville, TN 37235

douglas.h.fisher@vanderbilt.edu

Abstract

Region Radio is an artificially intelligent platform that finds and tells stories about the places that a user is moving through -- by car, bicycle, foot, or any mode of transport. Stories will typically be non-fiction with environmental and historic-cultural themes that are intended to increase a conservationist consciousness on the part of listeners and content contributors. Region Radio finds its stories on the Web, based on their relevance to places that lie on a route specified by the user; filters and ranks these stories on projected relevance; schedules these stories onto a playlist that corresponds to the route; and reads these stories as the user travels the route (to include “travel” in the minds-eye).

Introduction

Region “Radio” is an artificially intelligent platform that finds and tells stories, typically non-fiction, to people about the places that they are moving through -- by car, bicycle, foot, or any mode of transport. Region Radio finds its stories on the Web, based on relevance to locations that lie along a route specified by the user; filters and ranks these stories for projected relevance; schedules these stories onto a playlist of stories that corresponds to the route; and reads these stories as the user travels the route (to include “travel” in the minds-eye). “Place” is broadly construed to include the physical location, but also the people and events of the locale across time. Our current focus is conservationism, both environmental and historic-cultural, with the goal of increasing conservationist consciousness on the part of listeners and content contributors (Milligan, 2011).

We have developed an initial implementation of Region Radio, which as yet has not been thoroughly evaluated or deployed. The implementation is internally complete in terms of the most important basic functionality, though longer-term desiderata are not yet implemented.

Curation as Creation

The creations that Region Radio produces are playlists of existing stories that are found on the Web. The creative act is one of curating.

Outwardly, the playlist is a total ordering on selected stories, but should a user swipe past a story (not implemented)

associated with a place, another is selected from a location-sensitive priority queue of alternatives (implemented), and so the playlist is actually a partial ordering on stories. A playlist is a narrative, but one that is composed of loosely-coupled components (i.e., stories) as with most any curated collection. As narrative, we can assess the playlist in terms of characteristics like tension, polarity, emotion, and ultimately interestingness and user satisfaction. The latter characteristics have much to do with the characteristics of the individual stories, but also how the collection hangs together, which we discuss in Ongoing and Future Work. From a theoretical perspective, the **kind and amount of coupling** between components in a playlist are, perhaps, one extreme in a useful framework for studying creative narrative artifacts, where book chapters represent much greater coupling and reside in a different part of the proposed coupling spectrum/framework.

Region Radio has important connections to

- ♦ place-based education (Gruenewald & Smith, 2008);
- ♦ locative narrative (Greenspan, 2011);
- ♦ tour “guides” (Nisi, et al, 2008);
- ♦ points-of-interest notifications (e.g., Google Fieldtrip);
- ♦ configuration tasks (e.g., Maher, et al, 2016); and
- ♦ interactive story telling (Nisi, 2017), particularly in future designs.

Region Radio appears novel relative to other related efforts (in one or a combination of ways) in its *use of AI* to find place-relevant *stories*, with ambitions that these be *interesting* stories, from a *virtually limitless space*, and to *schedule these stories* into a playlist, with ambitions that the *interestingness of the playlist itself be greater than the sum of its parts* (i.e., the individual stories).

We believe that Region Radio will benefit three audiences.

Listeners

Region Radio is intended to educate users about places -- the geography, the environment, the people, the history, and the culture. The listener may be on a drive from Seattle, WA to Charlotte, NC. Alternatively, they may be on a frequently repeated walk/run in their community and its environs. Region Radio has enough intelligence to avoid telling the listener the same stories over and over, and

more generally, will customize a “podcast” to the user (but not yet) and the places in the regions that the user moves.

Authors

A distinct, and we hope overlapping population, are content developers -- “story tellers”. The stories that Region Radio finds on the Web can be from established sources (e.g., Nature Conservancy Magazine, community newspapers), but we view Region Radio as potentially incentivizing story authoring by community members of all types, including students, parents, grandparents, teachers, and churches. It is in this population of authors that we see the potentially transformative opportunities of Region Radio in place-based education. As an incentive mechanism, Region Radio adds to work on *spatial-digital story lines* or DSSLs (Hall, et. al., in press). Story authoring is a pipeline that includes collecting facts, memories, news, and insights, much like that in the map-based DSSLs. DSSL developers, and/or others, can translate these spatial narratives into text-based stories that can be found and told by Region Radio to the public on a potentially large scale. In the case of student authors of final textual stories, Region Radio can deliver their work to “authentic audiences” (Light, 2004), which may further incentivize them to produce accurate, comprehensive, and engaging material.

Other Curators

When Region Radio draws from, and acknowledges (as it always will), sources like magazines and newsletters, it supports these other curators and the cultural heritage that those curators help to maintain. As Madison (2011) notes “*analyses of creativity and innovation usually focus on producing new knowledge and offering access to it. Equivalent questions concerning existing knowledge, preserving and conserving old things and offering access to them, get less frequent attention.*” Because Region Radio draws from many other curators, it is intended to make listeners aware of these other curators, and to facilitate contributions to those other sources (see Ongoing & Future Work).

We turn to a system description of Region Radio, followed by an example playlist, ending with a discussion on basic research issues, and planned extensions.

System Description

Our initial implementation of Region Radio has been developed over 4+ months, distributed over the past two summers, by undergraduate researchers and developers. The implementation is in Python on a Github repository.

Figure 1 gives a high-level architectural diagram of Region Radio as it is currently implemented. Given a route to be travelled, the system constructs a playlist of stories about various places along the route. Each story in the playlist has an expected duration, and travel times are estimated as well. The playlist is constructed so that a story’s telling will be timed to correspond to a time when the cor-

responding places are being approached. The four depth-2 processes are arranged as nested loops, as suggested by the nested boxes, so that finding a suitable location from Location Analysis is informed by the stories that are associated with places in the vicinity of that location (i.e., through retrieval, evaluation, scheduling).

In this initial implementation, the story playlist is created offline, before actual travel, for upload to an appropriate listening device. We have represented the playlist as a partially ordered plan to facilitate dynamic adjustments in response to user decisions during travel, but at this writing we have not implemented a dynamic user interface that can exploit the partial-ordering of stories through explicit user selection of stories or through dynamic GPS updates on actual position. The travel route is created by Google Maps Directions API. In short, we describe the offline construction of a static playlist of stories, given a travel route that is an input to the system.

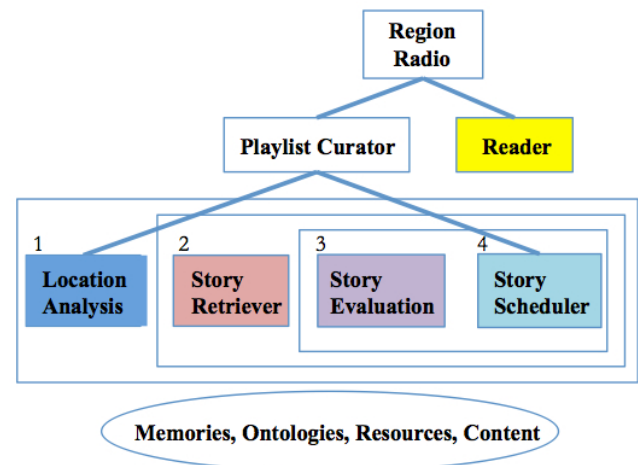


Figure 1: Region Radio high-level architecture

Location Analysis

In the outer loop of the Region Radio planner, the route is analyzed backwards from the goal destination. Place names within a radius of the goal location are identified using a Google Places API (e.g., of type ‘park’ or ‘natural feature’). The proximity radius used in this step is adjusted – if too few places are found in the initial circle, then the radius is increased; if too many place names are found in the initial radius, then the radius is decreased. After parsing through the JSON response to a Nearby Search, we also perform a Place Details Search on each location, to gather its coordinates, town, and state. Under the subsection “Closing the Loops” we describe how the next sub-goal location, again working backwards, is determined.

While we have experimented with semantic web based abilities to find terms that are related to the various place

names found above (e.g., who founded the place; what events occurred at the place), these are not yet integrated into our location analysis. Semantic relationships are found on sites such as Wikidata (e.g., Nielsen, 2013) and can be scraped from Wikipedia, and queries can be augmented with these terms to improve story retrieval.

Story Retrieval

As places are found, and in the future elaborated with semantic webs, stories are retrieved from the Web through a Google API using place-associated keywords. Rather than simply being called once with all available keywords, there is an internal, non-Web search through the space of keyword combinations, and the Google API is called for various subsets of keywords. In the future, semantic web connections will be used to guide the internal search of keyword subsets, but now the process is rudimentary, starting with a most general place name and adding other keywords if the previous search does not return stories of adequate quality. It is worth noting that the addition of even simple meta phrases to a search query (e.g., adding “story about” or “history” to a place name) can increase the density of interesting hits. Note that because there can be multiple place names associated with the goal location in the outer loop, there can be queries for each of these place names during the inner story retrieval loop.

Because we used the Google Custom Search API under a free account, our accesses were very limited in number. In our next phase of development (summer 2019) we plan to upgrade the account as necessary to allow and test more sophisticated and more numbers of Web searches.

Story Evaluation

While the stories are retrieved using place relevant keywords, this is insufficient to ensure place relevance of the results. So post-retrieval, Google Named Entity Recognition finds/tags nouns (e.g., categorized as “LOCATION” and tagged as “PROPER”). The Google geocoder is used to find the distance (e.g., Euclidean) between these locations and the query place. A hierarchical distance measure is also being investigated, but not integrated yet, using information from the Wikidata semantic web that links together places by a variety of relationships, such as inclusivity (e.g., Nantahala Forest surrounds Franklin, NC).

We have also investigated story “interestingness” at some length, but have not yet integrated it into evaluation. Story interestingness is critical so that we don’t end up reading Wikipedia pages, for example, as part of the playlist. Wikipedia and Wikidata are great sources of semantic web material, but often boring as stories. See Future Work for more on our story interestingness investigations. We have (so far) “sidestepped” the need for robust evaluation of story interestingness in Region Radio by biasing our stories to those that appear by existing curations that presumably vet for interestingness (e.g., Conservationist Magazine).

Story Scheduling

Stories from vetted sources that are deemed relevant to places in the vicinity of the (sub)goal location that is the focus of the outermost loop (under Location Analysis), are placed on a priority queue, currently based on relevance scores. The reading time of each story is estimated from the text length and an assumption of an English-speaker reading-aloud rate of 150 words per minute. The estimated reading time is used to estimate the time during the trip when the system must start reading the story aloud (using off the shelf text-to-speech) so that it ends just as the goal location is reached. Our assumption is that listeners would prefer to hear about places that they are approaching rather than places they have passed.

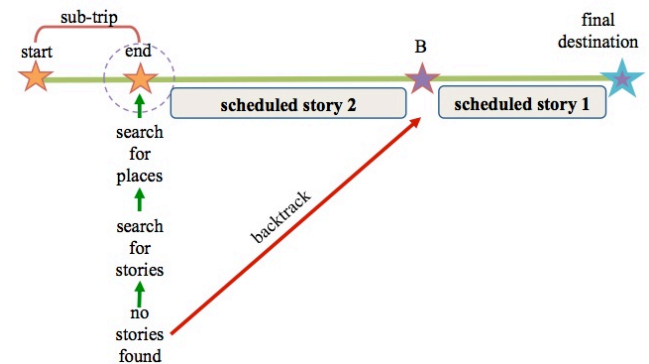


Figure 2: Backwards backtracking search for a playlist

The top rated story for the vicinity of the goal location is placed at the end of the playlist. The estimated starting point of this story becomes a sub-goal location to the planner, which next looks for stories that can be told as this sub-goal location is approached.

Closing the Loops

Figure 2 illustrates an intermediate step in planning of the playlist. The rightmost story was found first (using “final destination” as the goal), the story just to the left of it was found next (using sub-goal B), and its estimated starting point becomes a sub-goal for finding the next, third from the end, primary story in the playlist. Additionally, Figure 2 illustrates that if a third-from-end story cannot be found for the remaining unscheduled sub-trip, then backtracking occurs, and “scheduled story 2” will be retracted and another story drawn from the priority queue for sub-goal B. Backtracking search is a simple, elegant approach, which can be amended by procedures like keeping and tweaking the best schedules found so far in terms of unused time, which is particularly important to avoid repeated backtracking when unscheduled sub-trips become too short.

Other Current Functionality

Region Radio maintains a cache of previously found stories that can be reused as appropriate, currently to reduce

the number of Google Custom Search API calls. This cache will evolve into a repository that is shared across user accounts and be a store for a recommender system.

After retrieval, stories are scraped to eliminate non-story content (e.g., ads), to check for profanity (for eventual warning labels), and replacement of abbreviations (e.g., St. Mary's → Saint Mary's, rather than Street Mary's).

Stories are also introduced with original text that contextualizes the story, orienting the listener to the relevant places (e.g., "fifty miles northeast along US highway 64 is Franklin, North Carolina"). Eventually other contextual and factual information about the place, as can be found on Wikipedia, will be included before launching into a story.

After creation, the playlist stories are read aloud. We have experimented with free text-to-speech readers (e.g., gTTS, IBM Watson, Amazon Polly), with attention to intonation and an ability to adjust reading speed.

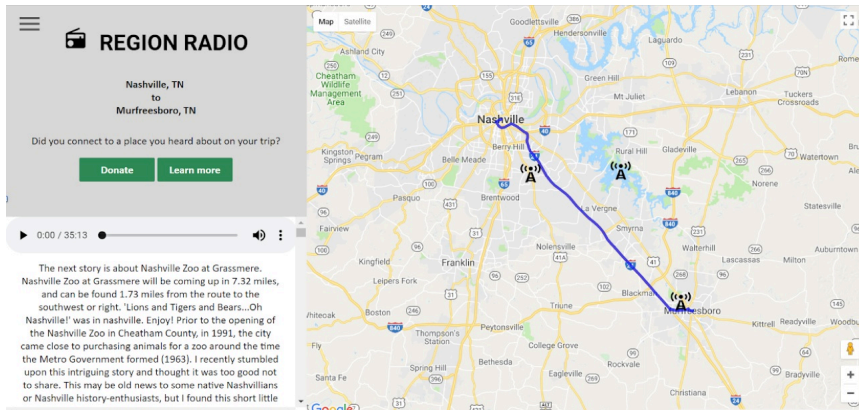


Figure 3: Web-interface to Region Radio

Figure 3 shows the Web interface to the current version of Region Radio, and a sample playlist. A route is shown, along with broadcast tower icons next to the places along the route that are the subject of stories. The first story to be read is about the history of the Nashville Zoo. Two buttons are shown. "Learn more" is functional and links to source files and other related material. The "Donate" button is not yet functional, but is explained shortly.

Ongoing and Future Work

We have mentioned some ongoing and future work (e.g., augmenting queries with semantic web relationships). We are also interested in **evaluating Region Radio** in terms of user satisfaction and whether it increases conservation awareness and place attachment. Long-term effects are hard to evaluate, but we are implementing the potential for donations to the source of stories (e.g., Conservationist Magazine) and the subject of stories (e.g., Joyce Kilmer Woods, Jefferson Street Sound). Tracking donations would be a quantitative measure of short-term effects on listeners, and donations might be incentivizing for curators to make their content available to us. The potential for donations

comes with potentially ethical issues, such as guarding against donations to some (e.g., hate groups). In some cases legal agreements with curators are necessary before using them in a public release of Region Radio.

User accounts/profiles are important for recording stories already heard, preference behaviors (e.g. preferences for wildlife versus energy stories), and eventually a support for recommender systems.

The stories in a playlist are currently coupled only by location in the playlist and implicitly by the proximity of the places the stories reference. We are interested in **thematic playlists** (e.g., Civil War sites, Civil Rights sites) using semantic information as mentioned earlier. We have investigated other **measures of story and playlist interestingness** (Ganguly, et al, 2014), which include characterizing narrative trajectories of topics and of sentiments. A playlist that covers distinct topics (e.g., wildlife to human settlement to carbon capture) and sentiments (e.g., habitat destruction to community development) may be more interesting than "flat" trajectories. We've also investigated trajectories within individual stories with sliding windows across the text. Characterizing interestingness is an important functionality that we continue to focus on.

We have investigated **text summarization** (Allahyari, et al, 2017) as a way of previewing and reminding users of story content. In particular, we will integrate story summarizations of one or two lines, which we have satisfactory results for, into the prefacing text to stories in the playlist.

Finally, we will integrate **dynamic capabilities** into Region Radio. Already underway is an allowance for users to swipe past stories, interpreting such actions as preferences, making use of alternate stories in story priority queues, and re-planning the playlist as necessary. Other dynamic capabilities like tracking trip progress and adapting the playlist in response to changes are longer term, after we move to mobile apps.

Acknowledgements

Supported by NSF #1521672 "Collaborative Research: CompSustNet: Expanding the Horizons of Computational Sustainability" and NSF #1623690 "EXP: Bridging Learning in Urban Extended Spaces (BLUES) 2.0." Special thanks to earlier Region Radio development team members, Mateus Winkelmann, John Kim, and Hannah Braun, and to members of the Space, Learning & Mobility Lab. We thank anonymous reviewers for their comments, to include calls for a framework of creative works in which curation might be situated, and for raising the potential for traumatic stories and otherwise out-of-place stories.

References

- Allahyari, M., Pouriyeh, S., Assefi, M., Safaei, S., Trippe, E. D., Gutierrez, J.B., & Kochut, K. (2017). "Text Summarization Techniques: A Brief Survey." In International Journal of Advanced Computer Science and Applications(IJACSA), Vol. 8 No. 10, 2017
- Ganguly, D., Leveling, J., Jones, G. J. F. (2014). "Automatic prediction of text aesthetics and interestingness." In Proceedings of the International Conference on Computational Linguistics.
- Greenspan, B. (2011). "The New Place of Reading: Locative Media and the Future of Narrative." In Digital Humanities Quarterly, Vol. 5 No. 3
- Gruenewald, D.A. & Smith, G. A., Eds. (2008). *Place-Based Education in the Global Age: Local Diversity*. Lawrence Erlbaum.
- Hall, R., Shapiro, B. R., Hostetler, A., Collins, H., Owens, D., Daw, C., & Fisher, D. (in press). "Here-and-Then: Learning by Making Places with Digital Spatial Story Lines." To appear in Cognition & Instruction.
- Light, R. (2001) *Making the Most of College: Students Speak Their Minds*. Harvard University Press.
- Madison, M. J. (2011). *Knowledge Curation*. Legal Studies Research Paper Series, Working Paper No. 2011-13. University of Pittsburgh School of Law.
- Maher, M.L., Lee, L., Gero, J. S., Yu, R., Clausner, T., (2016). Characterizing tangible interaction during a creative combination task. Proceedings of Design Computing and Cognition.
- Milligan, M. L. (2011). "Interactional Past and Potential: The Social Construction of Place Attachment." Symbolic Interaction, Vol 21, No 1: pp 1-33.
- Nielsen, F. Å. (2016). "Literature, Geolocation and Wikidata" Tenth International AAAI Conference on Web and Social Media, AAAI.
- Nisi, V. (2017). "The changing panorama of interactive storytelling: a review from locative to transmedia." DOC online - Revista Digital de Cinema Documentário. 43-68.
- Nisi, V.; Oakley, I. & Haahr, M. (2008). Places, location-aware multimedia stories: turning spaces into places. Artech, International Conference on Digital Arts. pp. 72-82.
- Weiss, A. S. (2013). "Exploring News Apps and Location-Based Services on the Smartphone." Journalism & Mass Communications Quarterly, Vol. 90, No. 3. Pp. 435-456.