

# ADA – Automated Data Architecture

## Creating User Journeys Through Content Using Linked Data

Jo Kent<sup>(✉)</sup>

BBC Radio & Music Multiplatform, London, UK  
jo.kent@bbc.co.uk

**Abstract.** The BBC has a wealth of permanently available programmes across a wide range of subjects with very low usage. We wanted to create a route into these programmes which balanced the need for curated, high quality journeys between programmes and the limited resource available for that curation effort. I will demonstrate ADA, a system created to create consistent, meaningful high-quality links between programmes with limited user input.

## 1 Introduction

There is a need for content providers to create consistent, high quality onward journeys to available content. Across the industry solutions used range from the heavily internally manually curated approach of Netflix<sup>1</sup>, to the user-driven algorithmically determined approach of Spotify<sup>2</sup>.

In this demo I will demonstrate ADA (Automated Data Architecture) which uses minimal manual curation and linked data to provide high quality serendipitous onward journeys.

## 2 Understanding the Problem Space

### 2.1 Assigning Metadata

The BBC has at least 34,000 permanently available speech radio programmes to which the traffic is low. There is no easy path into all available content. Some programmes have archive navigation, but these are isolated, specialised and heavily curated, and with decreasing team sizes, even these are not sustainable in the long term.

Our news and sport teams have long used linked data to dynamically populate article pages, which are set up using a strict, pre-existing ontology<sup>3</sup>. This constrains the browser

---

<sup>1</sup> Netflix manual tagging process: <http://www.techradar.com/news/television/netflix-wants-to-pay-you-to-watch-shows-here-s-why-1256098>.

<sup>2</sup> Spotify's algorithm explained: <http://qz.com/571007/the-magic-that-makes-spotifys-disc-over-weekly-playlists-so-damn-good/>.

<sup>3</sup> Sport ontology: <http://www.bbc.co.uk/ontologies/sport>.

to a rigid structure which may not match their world view. In any case, such an ontology does not exist across all programmes, the subject matter is too diverse.

Crowdsourcing metadata creation has been used by our R&D department on the World Service archive, this achieved at best 30.3 % precision (36.7 % recall)<sup>4</sup> largely because every person may have a different perception of the subject matter or the meaning of a term [1]. Also it was inconsistent: people only added tags to programmes which interested them, so some programmes have lots of tags and some none at all.

A team of researchers would provide better quality, more consistent data [2], but at a permanently high cost, in staff time. This is unfeasible with fewer staff available.

A fully automated system would not be able to deliver consistent quality standards for an audience facing offer: the automated interpretation of a homograph can give an erroneous or even offensive connection (e.g. Georgia the country as opposed to the American state). Any loss of data quality can cause a loss of trust in our content [3].

A middle ground needs to be found between these levels of automation, without compromising quality. We cannot expect producers to classify consistently, but they do know the precise subject of their programme. Therefore we need a system which only requires them to enter that subject, (e.g. a programme on autism can just be tagged with autism) without the need to classify the concept. Without this classification therefore, we need a system that will automatically supply the links.

## 2.2 Classification Systems

The ideal classification system would need to be recognisable and therefore trusted by our audiences and also be flexible and maintainable over time, as perceptions change [4].

Maintenance of our own ontology requires a significant staff time overhead, but the use of eternally maintained ontologies means we cannot control when changes happen, and still have to adapt when they do. Given the diversity of subject matter (subjects include the A470 (a road in Wales), Munch's "The Scream", virtue, Canada geese, existentialism and the Battle of Bosworth Field) the task of creating an ontology to cover and group every possible subject would be unfeasibly large. To make it manageable, we would have to make arbitrary choices about classification to make the multidimensional world fit in a two dimensional hierarchical structure. This is increasingly viewed as an outmoded and dictatorial organisational method, compared to open ontologies and collaborative folksonomies [5], and any arbitrary divisions of data are no longer semantic distinctions but simply an organisational tool.

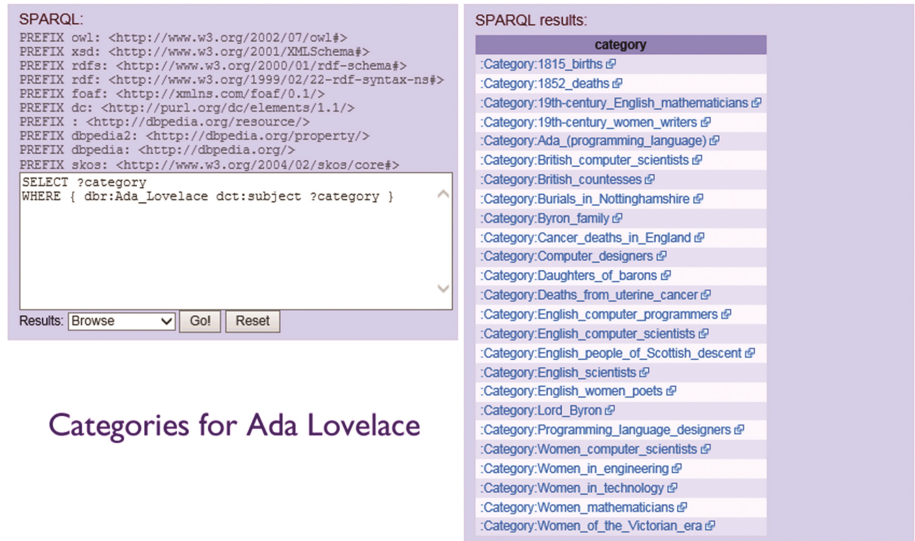
## 3 Unlocking the Power of Linked Data to Provide Automated Onward Journeys

The most promising linked open data sources were the Wikipedia/Dbpedia and Wikidata datasets. We found that there was no consistent hierarchical navigation or grouping

---

<sup>4</sup> <http://www.bbc.co.uk/rd/blog/2014/08/data-generated-by-the-world-service-archive-experiment-draft>.

information applied to the datasets. Wikidata has classification such as Library of Congress and Dewey Decimal mappings, but these are inconsistently applied<sup>5</sup>. It also offers classes and subclasses<sup>6</sup>, again inconsistently applied, which often simply cut off without reaching the top class of ‘Thing’. Dbpedia has classes<sup>7</sup>, but again these are only applied to a fraction of instances<sup>8</sup>. Dbpedia has categories for every subject, which offer a skos:broaden<sup>9</sup> journey to other categories, however, due to the way it is structured, often the category that was two hops broader was the initial category we started with, which meant that we had simply introduced more categories without any additional clarity (Fig. 1).<sup>10</sup>



Categories for Ada Lovelace

Fig. 1. Categories in dbpedia

Having found no usable hierarchical or grouping information we looked again at categories in Wikipedia/dbpedia. These have been added by Wikipedia editors, each adding the facts they felt were most salient. Anyone can remove them if they disagree, so they are effectively crowdsourced and peer reviewed. This means they have the recognisable relevance that people will respond to, while being of a high quality.

Asking producers to simply identify the subject for their programme means we can assure the quality of the initial reference, and automatically link to all of the categories (an average of seven per subject), which are matched to others to create user journeys

<sup>5</sup> Fewer than 3,500 Wikidata entities have Dewey Classifications attached.

<sup>6</sup> <https://www.wikidata.org/wiki/Property:P279>.

<sup>7</sup> <http://mappings.dbpedia.org/server/ontology/classes/>.

<sup>8</sup> Only 182 of 720 (25 %) sampled had types applied.

<sup>9</sup> <http://www.w3.org/2009/08/skos-reference/skos.html#broaden>.

<sup>10</sup> Using the 3.9 dataset, this issue seems to have been improved in the 1/4/2016 release.

that we could not create using manual curation without hours of research. At best a curatorial team might have added tags to Ada Lovelace like ‘computer scientist’ or ‘mathematician’ but here we have links to such diverse groups as ‘programming language designers’ and ‘British countesses’. These small, precise categories give a serendipitous feel to the journey and allow the users to learn more about the subjects even as they are navigating between programmes.

By discarding the notion of a hierarchy and instead presenting a graph, the journeys are not constrained in to a single worldview. We know from Lobel and Sadler’s work on homophily (i.e., love of the same) that “In a relatively sparse network, diverse preferences present a clear barrier to information transmission. In contrast, in a dense network, preference diversity is beneficial” [6]. So we can see that people respond better to a wide range of links that may not match their world view than to a narrow one. Therefore providing a broad range of linking categories (which have been selected by their peers as relevant) to each subject will present links the user will instinctively have a positive response to (Fig. 2).

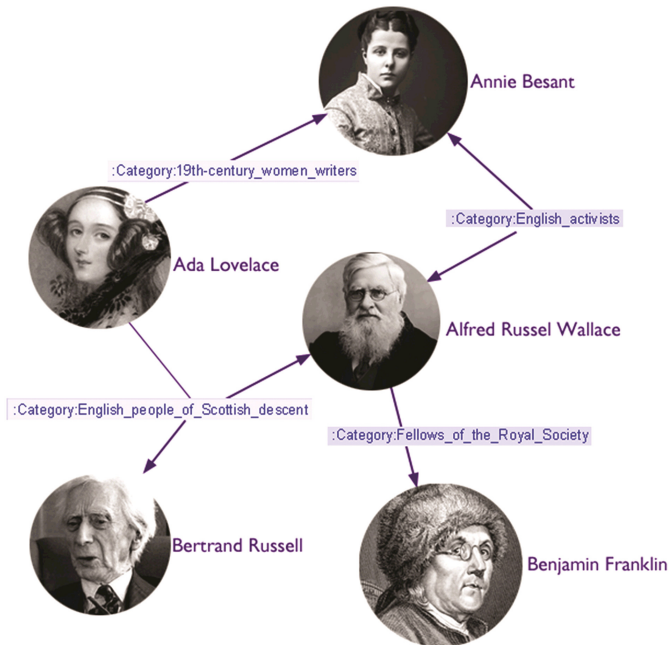


Fig. 2. Category links between people

## 4 Evaluation

Beginning with our initial sample of 610 programmes, we extracted over 1000 categories, of which 554 were linked to more than one programme, some of them to as many as 12. We only use categories which link to two programmes or more because only those

offer an onward journey. We keep the non-matching categories in the ADA triple store so that they can be used as soon as a new programme with a matching category is added. Some maintenance categories such as “World Digital Library related”<sup>11</sup> or “Articles with inconsistent citation formats”<sup>12</sup> were added to a blacklist as these are not useful user journeys. We were then able to examine the quality of the journeys offered. A programme on Roman Satire yielded links to 14 other programmes through five different categories; a greater and more detailed level of linking than in our bespoke archives.

We launched a beta<sup>13</sup> to gauge the audience reaction to the new navigation, and the response has been overwhelmingly positive with a rating of 4.15 (out of 5) stars on BBC Taster<sup>14</sup> and 164 (out of 250) positive verbatim responses through the demo feedback link. Once we rolled out ADA to all of our programmes, we plan to roll this out to other departments in the BBC to bring in news articles and educational literature and then to partner agencies (particularly cultural heritage organisations and learning institutions) to create learning journeys across all of our content by subject, rather than content type.

## 5 The Demo

In the demo I’ll be showing the beta and the API calls that power it. Visitors will be able to see and experiment with the semantic onward user journeys made possible by the use of linked data.

## References

1. Feyisetan, O., Luczak-Roesch, M., Simperl, E., Tinati, R., Shadbolt, N.: Towards hybrid NER: a study of content and crowdsourcing-related performance factors. In: Gandon, F., Sabou, M., Sack, H., d’Amato, C., Cudré-Mauroux, P., Zimmermann, A. (eds.) ESWC 2015. LNCS, vol. 9088, pp. 525–540. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-18818-8\\_32](https://doi.org/10.1007/978-3-319-18818-8_32)
2. Lykke, M., Lund, H., Skov, M.: Metadata in CHAOS: how researchers tag and annotate radio broadcasts. In: Knowledge Organization – making a difference (ISKO UK biennial conference) (2015)
3. Johnson, F., Saffi, L., Rowley, J.: Assessing trustworthiness in digital information. In: International Data and Information Management Conference (IDIMC) (2014)
4. Cupar, D.: Diachronic semantics: changes of meaning of words over time and the consequences for keeping classification systems up to date. In: Knowledge Organization – making a difference (ISKO UK biennial conference) (2015)
5. Busch, J.: Web-based content organization and the transformation of traditional classification systems. In: Knowledge Organization – making a difference (ISKO UK biennial conference) (2015)
6. Lobel, I., Sadler, E.: Preferences, homophily, and social learning. In: Operations Research (2014)

<sup>11</sup> [https://en.wikipedia.org/wiki/Category:World\\_Digital\\_Library\\_related](https://en.wikipedia.org/wiki/Category:World_Digital_Library_related).

<sup>12</sup> [https://en.wikipedia.org/wiki/Category:Articles\\_with\\_inconsistent\\_citation\\_formats](https://en.wikipedia.org/wiki/Category:Articles_with_inconsistent_citation_formats).

<sup>13</sup> <http://www.bbc.co.uk/inourtimeprototype>.

<sup>14</sup> Feedback on BBC Taster overall rating 4.15 stars: <http://www.bbc.co.uk/taster/projects/in-our-time>.