Drag it together with Groupie: Making RDF data authoring easy and fun for anyone

ABSTRACT

One of the foremost challenges towards realizing a "Read-write Web of Data" [3] is making it possible for everyday computer users to easily find, manipulate, create, and publish data back to the Web so that it can be made available for others to use. However, many aspects of Linked Data make authoring and manipulation difficult for "normal" (ie non-coder) end-users. First, data can be high-dimensional, having arbitrary many properties per "instance", and interlinked to arbitrary many other instances in a many different ways. Second, collections of Linked Data tend to be vastly more heterogeneous than in typical structured databases, where instances are kept in uniform collections (e.g., database tables). Third, while highly flexible, the problem of having all structures reduced as a graph is verbosity: even simple structures can appear complex. Finally, many of the concepts involved in linked data authoring - for example, terms used to define ontologies are highly abstract and foreign to regular citizen-users.

To counter this complexity we have devised a drag-and-drop direct manipulation interface that makes authoring Linked Data easy, fun, and accessible to a wide audience. Groupie allows users to author data simply by dragging blobs representing entities into other entities to compose relationships, establishing one relational link at a time. Since the underlying representation is RDF, Groupie facilitates the inclusion of references to entities and properties defined elsewhere on the Web through integration with popular Linked Data indexing services. Finally, to make it easy for new users to build upon others' work, Groupie provides a communal space where all data sets created by users can be shared, cloned and modified, allowing individual users to help each other model complex domains thereby leveraging collective intelligence.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General terms: Design, Human Factors Experimentation

Keywords: Guides, instructions, formatting.

INTRODUCTION

In this paper we describe Groupie, an experimental UI which aims to allow causal users to create and edit structured relational data directly on the web, something that to date has been the purview of expert users and expert tools alone. As we discuss in the background section, not only do we find this an interesting software and interaction design challenge, we see several reasons why such a tool is increasingly necessary for people to be able to interact with and make use of the ever-increasing quantities of data being published on the web.

The data we are describing is Semantic Web data, or what is more currently becoming known as Linked Data [4]. Linked data is most often expressed in RDF [2], the Resource Description Framework language designed to be able to support sharing and representation of a large class of structured data. There are several reasons why RDF is a useful format beyond, say, the spreadsheet, CSV, or XML for representing such data. Understanding these speaks to the challenges of the approach we propose.

RDF creates schema-free data, meaning that any data can be defined with an arbitrary set of attributes. The fundamental representation of RDF is a triple, where data is encoded as a relation in terms of a Subject, a Predicate, and an Object. Each component of a triple is identified by a universal identifier known as an URI, which, in turn can be "dereferenced" to map to a concept that provides more detail about what it represents. Thus, unlike XML, any node in the graph is specifically addressable. These addressable relations enable possible linkages between heterogeneous datasets to be discovered, and those connections to be leveraged in order to link several data sets together. For instance, If one data set lists a university with an address and address has street and has "pcode," and another data set lists library names but the only address component is a "postal_code", both sets URIs for the pcode/postal_code triples may either explicitly reference the official Postal Code data created by the post office, or can be mapped to that data set, which also translates postal codes to latitude and longitude coordinates. Thus in this case, these data sets can be plotted onto the same map to see what the distances are between libraries and universities without putting all the data into a single database (or mashup) and querying over that.

With these features have so far come costs in terms of how easy it is to create web data. To date, existing web data authoring and publishing tools have assumed that people creating such linkable data are both familiar with the RDF language, with URIs, and with how to define ontologies to represent as classes, types and properties the concepts they then

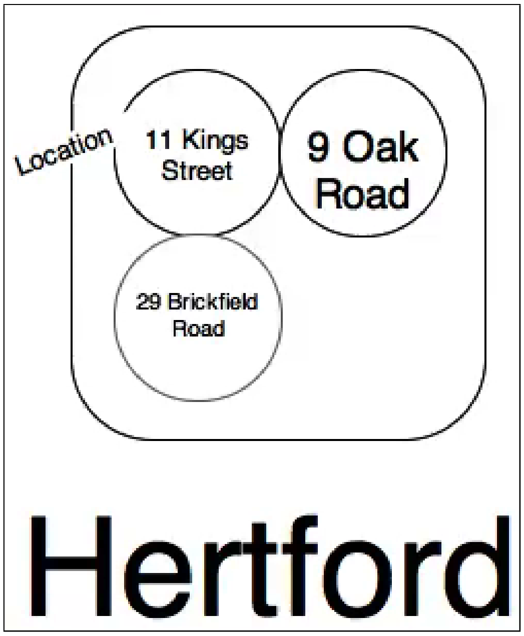

Figure 1: Groupie's representation that three addresses (circle objects) have a location (grouping object) that is Hertford (large text)

wish to populate with data, eg they will define what a library is in relation to other concepts before adding data to those attributes like its catalogues, addresses, staff and so on.

To make web data authoring more accessible, our approach has been to explore how much of this formal knowledge about web semantics we can obfuscate to empower a person to be able to create these kinds of semantically rich data sets. The result is Groupie. Groupie hides just about all references to RDF, ontologies and URIs from the user in order to present people with a more conceptually simple, direct manipulation interface that foregrounds one main concept: defining groups and putting things in groups. In Groupie, a large text label represents an entity; a circle represents entity attributes; a rounded rectangle represents a group of attributes that are part of an entity.

Groupie's representation of text, rectangles and circles fundamentally breaks with the formality of other RDF editor approaches that use a form creation/form filling model. By using objects instead of forms, our hypothesis has been that, like mind mapping or outlining tools, object manipulation makes it easier for people to jot ideas down for their data representations and to refine these relations as they developed them. This grouped view, as we describe below, also provides a visualization advantage over standard network graph representations of RDF data [11] to help a data author make sense of their information space. Last, our implementation makes it easy to reference or simply clone other people's models of data so that one may start from complete scratch or partial scratch to begin to develop linkable data.

We call Groupie our "lab-rat UI" because our main concern has been to explore how far back we can push the complexity of RDF/URIs/Ontologies to produce a lightweight UI that a non-expert can use and still create valid, usable RDF.

In the following sections we describe our motivation behind making RDF generation more accessible for regular web users. We present a detailed walkthrough of the tool, a discussion of the related work, and the results of an exploratory study to calibrate the performance of Groupie against the standard RDF tool. We conclude with next steps for the research stemming from this work.

BACKGROUND AND MOTIVATION

The focus of Groupie is to let anyone who might write a blog post or create a tweet also create open linkable data. The web as we know it right now, is still mainly a document web. Most individuals create microblogs - or tweets, facebook entries or blog posts. We may tag other entities, like friends in photos. The tools we have as users to engage with the Web are largely text or video upload based.

Hypercard Inspired - Casual Data Authoring. Once upon a time, though, user-oriented, data (not document) creation applications, like Douglas Englebart's NLS [5], fired the imagination of the nascent User Interaction community. Indeed, HyperCard [7], an early Apple Mac application traces its lineage to NLS. HyperCard let anyone who could point and click create a Stack of data (not primarily document) Cards. Data elements within the cards from characters in plays to ingredients in recipes could be interlinked for hyper-navigation within a stack.

HyperCard stack creation informed early multimedia tools – such as those by Macromedia for the authoring of interactive encyclopedias, books and games; stacks made by individuals from bibliographies to movie indexes were shared on bulletin board services and became some of the first ways the Humanities began to explore the potential for computers in what has become Humanities Computing. [15]

HyperCard stacks had several key qualities for creating data: a simple interaction model, form/field creation, and interlinking within cards in stacks. While hypercard itself had a full fledged language that enabled no small control of stacks as well as its operating system environment, entry-level users could still create cards and stacks via its simple UI in a tool that privileged interlinked data creation.

Despite all the power and greatness of the Web, we have lost most of this ready access to generating personal, portable data-oriented resources in favor of page-oriented resources. True, most of the commercial web is driven by content management systems linked with databases, but for the individual blogger/tweeter/facebooker/searcher, the current page-oriented or document web does not lend itself to casual data generation by lay users. Stacks, alas, are gone. With the explosion of scientific data, and especially government data being published to sites like data.gov and data.gov.uk, we don't believe the drive to create data rather than documents has necessarily gone away with the hypercard stack; It's just moved out of reach for most of us.

With Groupie, we wish to begin to explore if we can address the complexity of web data generation by offering an interaction with the directness of HyperCard for representing relations, but augmented with the interlinking, collaborative potential that is part of the Semantic Web to enable discrete datasets to interconnect. In other words, if I created a stack on HyperCard of guppies that talked about their color and size, and you created one about their feeding and life spans, it would have been very difficult to blend them into a third stack that let us explore all four attributes in one stack. With RDF generation, it is possible to generate these two sets as separate "stacks" of data but also blend them for just such views.

Locked Structured Data A second motivation for Groupie is to help create more open personal, social and public data. Right now, if we consider just the domain of personal or social information, copious amounts of social data both exists and is generated by users every minute, but most of it is currently locked into applications like Facebook, Twitter and Last.FM. This means that the opportunities to create new visualizations or applications that could leverage this data are at best difficult to construct, relying on the arbitrarily changing terms and conditions of these corporations' API and data policies. Even fabulous Web 2.0 mashups of Google Maps with any location-based data sets exemplify the problem: data for such mash ups has often been scraped from other sites, and is then itself hidden behind the wall of the application. As a fixed application as well, no new data sources can be added to it for arbitrarily new kinds of queries. A web 2 mash up that shows only the top ten universities mapped does not allow for a new data set about crime type by region to be added to it so that one can see which university is in the safest neighborhood.

A motivation for Groupie therefore is to see if a HyperCard-inspired approach to Semantic Web data creation will itself inspire people to create inter-linkable reusable open data the way people once created and shared HyperCard stacks. The work described in this paper is a first step towards this larger experiment: before people could create stacks, they needed HyperCard. Before the same kind of people who created stacks can easily create RDF, they too need a tool. We propose Groupie as a first step to test the hypothesis that usable RDF can be generated via a simple direct-manipulation GUI that hides under the UI the complexity of RDF-generation. The following section overviews the Groupie approach.

INTERFACE WALKTHROUGH

To provide an overview of how Groupie works, we begin with a simple scenario of how a person might create a model to keep track of real estate they have viewed and are considering purchasing.

When our user, Jackie, opens Groupie, she is first presented with a list of all the models she created previously, as well as worlds from other users which have been shared publicly. These worlds can be searched and browsed by author, date, and complexity.

When she sees a world about real estate in her area, she clicks 'Duplicate' to make a copy for her own editing. Several houses already exist in this world and she wishes to add a new one she has recently viewed at 29 Brickfield Road. The user starts in the 'All entities' view which shows all the resources in the world.

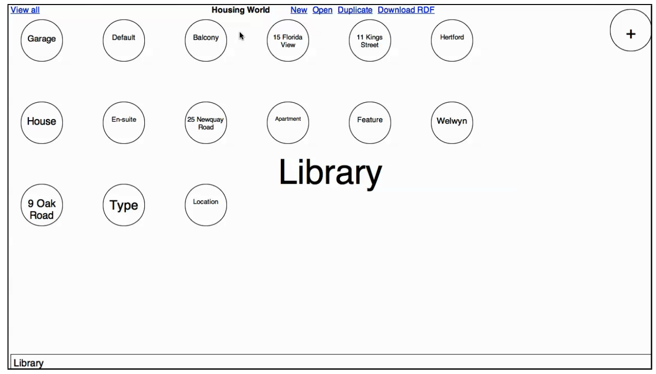

Figure 2: The 'all entities' view

The new house is located in the town of Hertford, so she double-clicks the entity representing Hertford to bring it into the view. The Hertford view shows everything known about Hertford - currently the fact that two houses - 11 Kings Street and 9 Oak Road are located in it.

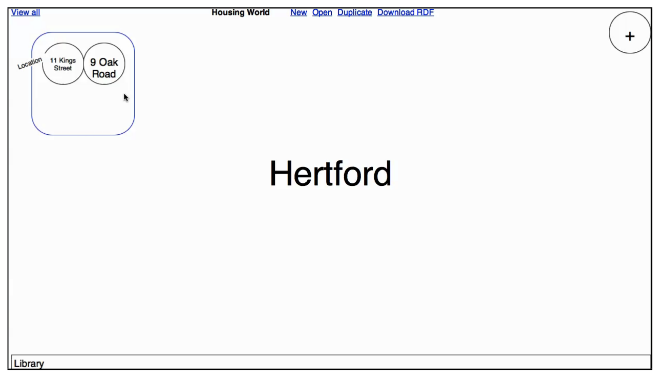

Figure 3: 'Hertford' view

To create a new entity, Jackie drags the circular '+' from the top right corner into the 'Location' group. This results in a new, as yet unnamed, entity which already has the property Location = Hertford set on it.

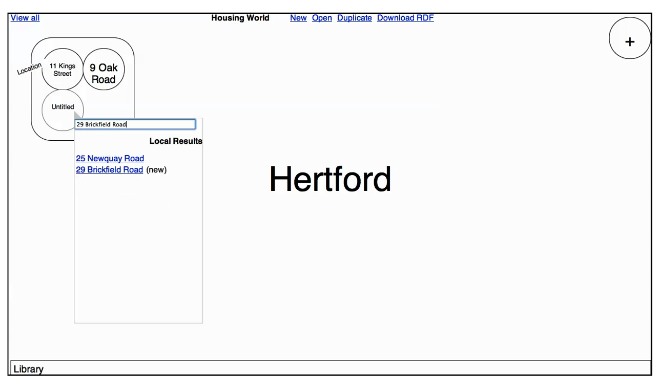

Figure 4: Naming the entity

The aligner window appears which allows Jackie to give the entity a name - in this case "29 Brickfield Road".

Next she double clicks on the entity to it bring into the main view and begin making more assertions about it. To say that the house has a balcony the user drags the + entity into a ‘ghost group’ which is an empty group which appears during all dragging operations and facilitates the creation of new groups.

Figure 5: Using the ghost group

Once again, the aligner appears. As Jackie types ‘Balcony’ the aligner performs a typeahead search which shows an existing ‘Balcony’ entity has been added previously. The user selects this. By aligning to existing entities rather than creating new ones, the value of the data set is increased by removing fragmentation and duplication of concepts.

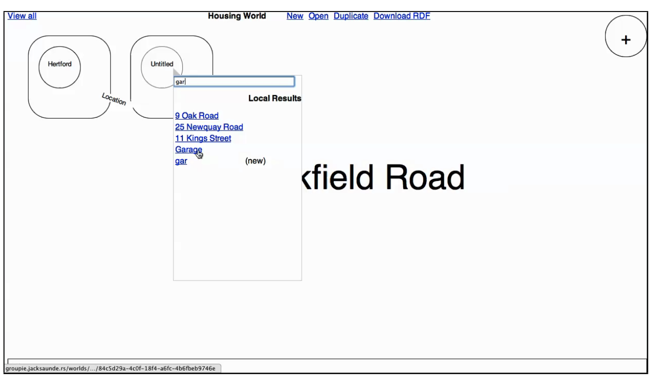

Figure 6: Aligning to ‘Garage’

An alternative technique for using an existing entity is to use the ‘library.’ The library provides quick access to all entities known about in the user’s worlds. To assert that the house has a garage, she simply picks up the garage entity from the library and places it into the same group as the balcony.

As this group does not yet have a name, Jackie needs to right click on it and choose ‘Rename.’ This presents the aligner window once again. The group can be aligned in the same way entities can.

Another group is created to represent the fact that this is an apartment as opposed to a house, by dragging the + widget into a new ghost group and aligning appropriately.

Now Jackie can view the data in several ways. In addition to viewing each individual house, they can be faceted into groups based on their location by bringing the ‘location’ entity into the main view by double clicking on it. This results in the same familiar grouping interface, but rather than showing properties containing values, it shows values containing subjects. This results in a faceted view showing houses grouped by location.

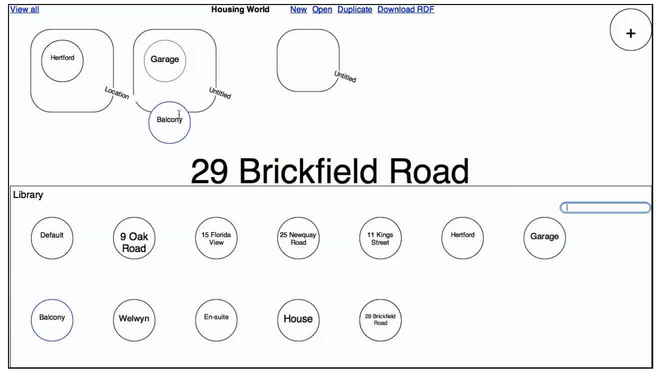

Figure 7: Using the library

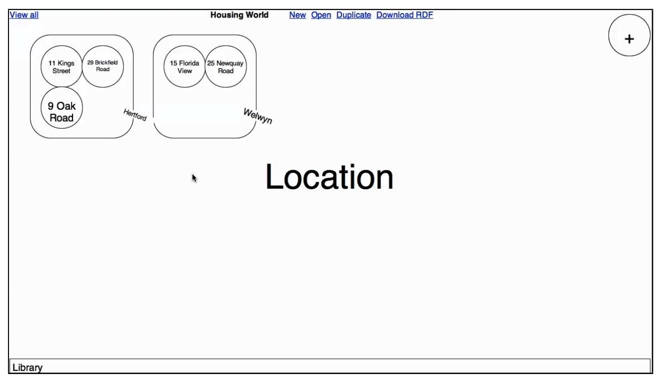

Figure 8: The ‘Location’ view

A similar view can be found by bringing ‘feature’ into the main view. This once again results in houses faceted by the features they contain - garage, en-suite and balcony. It is important to note that in this view, as in any other, the same entity can appear more than once. For example ‘11 Kings Street’ appears in both ‘Garage’ and ‘Balcony.’ To show their connection, all entities in other groups which are the same are shown highlighted when the mouse is moved over any one.

Finally, Jackie is able to publish and share the model she has created both to other Groupie users and on the wider semantic web as RDF, using the ‘Publish as RDF’ button, which results in a URI be provided which can be used with other RDF software.

RELATED WORK

Groupie draws its design and inspiration upon two types of tools– the first are tools for creativity support and external-cognition, including generating and organizing ideas, planning, problem-solving and sense making [16]. The other set are those for structured data creation and editing, including RDF linked-data authoring tools designed to facilitate the construction of RDF and OWL ontologies.

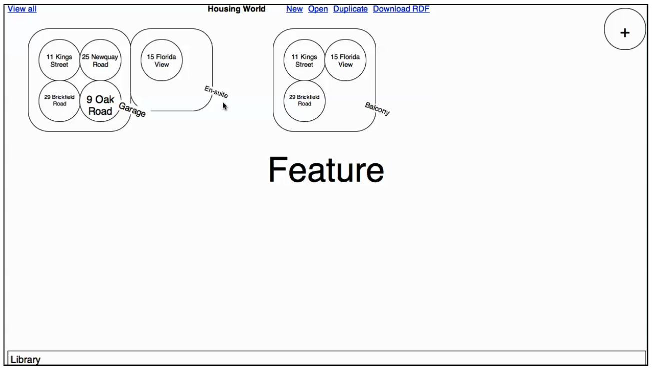

Figure 9: The 'Feature' view

Brainstorming and mind-mapping tools

Controlled studies comparing the use of mind-mapping tools to traditional approaches in such activities as brainstorming and studying have found a number of benefits of using these tools. Such benefits include improved later recall from memory [6] [14], more generated ideas, and more dense association connectivity in brainstorming sessions[12]. Thus, if these metrics are taken to reflect the effectiveness of these activities conducted with these tools, these tools be said to actually amplify these activities or at least to preserve more of their output for later use. The key affordances cited in these studies are that they helped people to articulate quickly and easily ideas and to freely make annotations among them. With similar affordances, our design approach has been that Groupie will well support such external cognitive activities and thus help enable and inspire equally rapid, light weight and exploratory data generation.

Towards encouraging creative collaborative exploration, *viki* [1], a colorful, visual wiki environment designed for kids used heterogeneous shapes and colors in a simple, drag-and-drop interface evoking the metaphor of a "sticker book" to permit the construction of linked collections of pages. Interaction in viki was restricted to the construction of new stickers, which could be created by dragging existing stickers created by the user or others onto a canvas. Such stickers were automatically shared with other users, and could be used to further derive new stickers. This simple and graphic Viki was found to be easy enough to be used by across a wide range of ages (children through teens), many of whom used the system to construct elaborate storybooks out of collections of sticker pages.

To make for a similarly light weight and potentially fun experience authoring RDF, Groupie's was kept visual and simple, and is likewise designed to be an inherently collaborative system – allowing one person to directly build upon others' worlds, and to refer to entities created in others' worlds to create inter-linked data sets.

Structured Data authoring tools

Most current linked data authoring tools target experienced linked-data knowledge engineers, providing advanced functionality such as reasoning, large-scale visualization, and interfaces with legacy databases. The drawback, however, of such functionality is complexity at the user interface; these interface expose nitty gritty details of the authoring process, exposing options and distinctions that may be not only irrelevant to the lay citizen user, but also confusing and potentially intimidating. Given the resulting learning curve, these tools are not designed for casual use.

Another design aspect of these tools is that they put ontology design first, allowing instances to be filled in only after the ontology has been defined. Yet, in working with experts to develop ontologies over the past five years, we have seen that people are inherently bad at devising ontologies upfront, and before instances. Increasingly, members of the Semantic Web community in particular are coming to the view in private conversations that perhaps the data first/ontology later view has value. That said, data generation tools are still for experts.

A mere handful of tools thus far have focused on structured data authoring for "Real" end-users. VITE [9], for example, is the most similar to Groupie in design in that it provides a direct-manipulation, shape-positioning metaphor to allow the authoring of structured data. In provided examples, users place shapes in a grid to express related concepts. But unlike groupie, the mappings between these visual and spatial attributes and their semantics have to be specified upfront by the user. While this customizability allows for much greater flexibility, it be unnecessary for casual users and novices, and could potentially serve as a use barrier that may deter casual users. VITE is also not designed specifically for RDF/linked data, providing no facilities for linking to external concepts or ontologies.

Other related tools

Exhibit [10] demonstrated that empowering individuals to easily publish structured data collections on the web revealed a significant need and desire to do so. Exhibit, however, focused on the publishing process, rather than the authoring/creation process; published collections were not editable.

DESIGN AND IMPLEMENTATION

In groupie we have taken on a number of issues that bridge the simple data representation and interaction model within groupie and the more complex aspects of creating actually usable, publishable RDF. We touch on several of these concerns as they influenced the design of Groupie.

Ontologizing before authoring As noted above, linked data authoring tools typically prioritize the creation and editing of ontologies, abstract descriptions of what the data will look like, or what components it will include (what are the properties of a friend? of a book?) before actually authoring data instances. The argument for upfront ontology construction is so that, once created, instances can be created and verified against this structure so that correctness and consistency can be guaranteed for data integrity and on the web, machine legibility

Without concrete examples of what they want to model to begin with, it can be difficult for people to ontologize from the outset: it can be a cart before the horse experience. If taken from the perspective of external cognition, in many such activities, the modeler may not be entirely sure what they are going to need to model, much less how they are going to model the material before starting the process. Such a process is typically begun by starting at one or several simple places, taking best guesses, and then successively and iteratively connecting, defining, and refining until a solution or representation is achieved [13]. Thus, in such a setting, by the time the ontology is known, the problem has already been solved.

We sought to design Groupie so that, like external cognition and mind-mapping tools, we do not put ontology design questions before instance design, but allow direct playful exploration and refinement of instances so these can be used encourage exploration of the space of potential representations without commitment, but once one is happy with the result, new ontologies may be designed from these examples, or appropriate references to existing ones accessed.

How simple can we make it?
In order to simplify the complexity of the linked data model, we started not with its smallest unit - a triple - but with a conceptual model's smallest unit, an entity. The approach is similar to conceptualizing the document web's smallest meaningful unit as a page rather than as it's smallest actual units, such as elements and attributes within a page like pictures, tables, text.

Putting related things together is a familiar action in daily life and in computer interfaces: food goes in the kitchen; shoes go in the closet. Likewise applications are generally found in an application directory; data, everywhere else. We tend to search for information in terms of groupings: by related terms, by dates created, by kind. Indeed everyone in our study was readily able to create entities by clicking the plus sign by the circle to make a new object and associate that object with a group.

We enrich what we know about something by linking it with other possibly related properties: a county has houses, CO_2 emissions, crimes, neighbors, hospitals. Each of the attributes of a county can itself be an entity made up of other groups. Hospitals have operating rooms, wards, morgues, nurses, drugs, budgets, patients. Patients likewise have diseases, treatments, and homes to which bills may be sent.

One of our main questions in evaluating Groupie however has been how well this simple interaction captures the model of both an entity and associating other objects within related entities.

How to achieve linkage with the cloud? One of the strengths of the Semantic Web is the opportunity for data re-use. At the top of the paper we described a library data set and a university data set making use of a third party postal code data set. By linking out to other data sets or ontologies where postal codes have been defined and/or instantiated, we don't have to reinvent the object; we can simply reference it to enrich our entities.

Marti Hearst in her discussion of tags and how their value increasing as tag use converges suggests that auto completing input mechanisms are effective at improving tag convergence in folksonomies [8] . In Groupie, we leverage this same convergence effect by using autocomplete predictive search at every moment the user needs to reference a resource or property – a search that encompasses both the resources defined in any of the user's own created worlds or in the greater linked data cloud via RDF discovery services like Sindice [17].

Visualization of the Graph One of the subtle properties of the Entity-by-Group-Creation aspect of Groupie UI is the view it offers on the data. In the Semantic Web, RDF or ontology data is often presented visually at best as a BFG (big fat networked graph [11]) where all triples have equal weight in a view. With its emphasis on entities rather than triples, Groupie foregrounds views of one entity (for example Hertford) and its associated groups (locations) and (features) at a time. Our sense so far is that such an entity-oriented view of one's model/data makes it easier to define and to explore relations within the information than does the traditional BFG.

Sharing and collaborative authoring One of the challenges of Semantic Web data authoring is collaboratively working on developing or revising data sets [2]. This challenge is understandably often about privileges: does one want budgetary information for hospitals to editable by the world? Open for comment? What about copies of the data? In Groupie, we employ a lightweight approach for shared editing or whole model cloning.

Groupie supports group authoring of information letting people easily share, build upon and copy others' work to help them get started. This done by providing a communal "gallery space" where the data sets (deemed micro-worlds) people create get automatically published, which can then be used to by other users to create their own micro-worlds. This sharing makes it easier to create RDF for two reasons: first, it's easier to pick up where someone else left off because new users can imitate what others have done instead of having to figure it out for themselves. Second, sharing worlds lets people distribute the work for modeling a complex domain. Finally, sharing worlds lets people more easily improve their data sets by letting them compare how they have done it with others, and borrow others' approaches.

Importing Existing Data
Once a user has chosen to align an entity to an existing RDF resource on the wider semantic web (via the autocompletion labeling), they may also wish to import some or all data relating to this resource. A naive approach to this would be to simply import all triples available in the resource, but this has several problems. Crucially, many RDF resources contain a very large number of triples, over 1000 in some cases, which is a scale too large for the groupie interface to be useful for. Additionally, users may not be interested in some triples or feel they are relevant to their purpose.

For these reasons, a property selector interface has been developed. This displays all triples found in a resource to be imported in a table view. Users may then pick and choose which triples to import into the Groupie interface.

STUDY
We used a combination of methods to examine the suitability of Groupie in a controlled study of several simple data au-

thoring tasks. This sections describes our methods, tasks and results.

Methodology
An A-B evaluation was performed to compare the usability of Groupie against existing RDF editor Protegé. 10 participants were recruited (ages 18-50) and given three tasks to perform. Each task was performed twice, once in Protegé and once in Groupie.

The first was a warm up task led by the investigator to provide instruction and training on how to use each interface. This task involved describing a picture of three wooden blocks. Observations from this task were not included in the study.

The two main tasks users performed are shown in Figure 10.

All interactions were recorded using screen recording software, and the produced models were saved for analysis. Finally, users were asked to complete a short questionnaire about the tasks afterwards.

Quantitative analysis To evaluate the correctness of models created by the users, we compared each of the models created by users against "ground truth" models constructed by an experienced linked data knowledge engineer. This knowledge engineer was instructed to keep such models minimal, and suitable for small-world web publishing (in order to avoid over-engineering problems frequently encountered in linked-data modeling), and to devise appropriate ontologies on the fly instead of only referencing existing ontologies. The ground truth model were refined by a second linked data engineer, who suggested improvements to the first. Then each user's result for each task was compared to these ground truth models by another expert along the following criteria:

Model quality metrics:

- *Entity correspondence* : Are each of the entities in the ground truth accounted for in the user's model? Do each of the entities in the user's model have a corresponding entity in the ground truth?
- *Relation correspondence* : Are each of the relations between entities in the ground truth accounted for in the user's model? Are each of the relations between entities in the user's model expressed in the ground truth?
- *Intra-model consistency/duplicate check* : Does each entity have only 1 representation in each model? Similarly, is each property expressed in one and only one way wherever it is used?

These metrics were chosen as intrinsic model quality metrics because they reflected how close each of the user's models arrived to the ideal models produced by an expert. The two correspondence metrics were computed as precision/recall measures from information retrieval, because, although this is in reverse from the typical information retrieval setting where the user is searching for information in a large database, the resulting elicited knowledge can be seen as consisting of representations (concepts and relations) that are relevant to the modeling exercise, and those that are not. Therefore we apply the metric of comparing the "correctly elicited relevant concepts" (true-positives) to all of the elicited concepts (True positives + False positives) to the concepts that should have been included but were not (False negatives). Since both the concepts and their relations are important to making a complete model, we compute such metrics for entities declared and properties and report on these separately.

The final intra-model consistency metric was chosen because it reflects the extremely common and challenging problem of co-reference failure that plagues linked data, in which several entities referring to the same extensional entity are referred to in an inconsistent manner.

The lack of standard metrics for evaluating the degree of external linkage led us to omit such metrics from this evaluation. Any such metrics would likely be sensitive to the kinds of things being modeled – e.g. domains that have already been modeled in well-known domains (such as publications, or biological concepts) would exhibit more external linkage simply because such domains were more commonly modeled, and had more established vocabularies than other domains. Since such considerations were not relevant to the quality of the authoring interface, we omitted such metrics.

In addition to the model quality metrics, we also measured the total time to completion for each of the modeling exercises.

Qualitative metrics
In addition to the model quality metrics, we captured all screen interactions which were analyzed after the experiment to permit the post-hoc comparison of how users used both interfaces. We also asked all participants to fill out an exit questionnaire, in which they were asked compare their experiences between the two systems, and for ways to improve Groupie.

Results
Model quality Results of the model quality analysis are available in Figure 11. As described earlier, we separately consider the entities created from the relations created among them.

The top row represents the ground truth model created by the linked data expert, and displays the number of entities and relations in both truth models. Each of the subsequent rows represents the responses of one user using either the Protégé or Groupie. Users used both interfaces for each exercise. The "# of entities/# of properties" columns represent the number of entities and relations created, respectively by a user in each interface condition. The "# of correct entities/properties" represents the number of entities that could be matched to an entity or relation in the ground truth model, as specified earlier. Precision, recall, and overall F1 metrics were then computed for each.

As can be seen, users produced high-quality models overall using both Protégé and Groupie for the Menu exercise (F1 for entities = 0.90, properties = 0.70), but experienced considerable difficulty with the Shakespeare exercise with both clients (F1 for entities = 0.63, properties = 0.55).

Comparing performance with each client, Protégé exhibited a slight advantage over Groupie with respect to precision in

Task 1

(No help to be given)

Groupie version: Go to
[illegible]rlds/09639dd6-6a7d-11e0-a7b5-1231380c15e5/duplicate

Protegé version: Open <Protegé data file provided>

Task:
Bob the Chef has prepared today's lunch menu. Add the following dishes to it:

Florentine Pizza tomato sauce, mozzarella, spinach and egg

Fish Cake with spinach and hand cut chips

Task 2

(No help to be given)

Familiarise yourself with this description:

House of Capulet

- **Capulet** is the patriarch of the house of Capulet.
- **Lady Capulet** is the matriarch of the house of Capulet.
- **Juliet** is the daughter of the Capulets, and is the play's female protagonist.
- **Tybalt** is a cousin of Juliet, and the nephew of Lady Capulet.

House of Montague

- **Montague** is the patriarch of the house of Montague.
- **Lady Montague** is the matriarch of the house of Montague.
- **Romeo** is the son of Montague and Lady Montague and the play's male protagonist.
- **Benvolio** is Romeo's cousin and best friend.

Create a new world/document and model either the Capulets or the Montagues.

18/4/11 version 1.1

Figure 10: Evaluation tasks given to users. (Warm-up task not shown.)

MENU

	condition	# entities	correct entities	entity precision	entity recall	entity F1	# props	correct properties	property precision	property recall	property f1
	Gold	6.00					8.00				
User 1	groupie	7.00	6.00	0.86	1.00	0.92	8.00	7.00	0.88	0.88	0.88
	protégé	5.00	5.00	1.00	0.83	0.91	6.00	6.00	1.00	0.75	0.86
User 2	groupie	6.00	6.00	1.00	1.00	1.00	5.00	4.00	0.80	0.50	0.62
	protégé	6.00	6.00	1.00	1.00	1.00	6.00	2.00	0.33	0.25	0.29
User 3	groupie	6.00	6.00	1.00	1.00	1.00	7.00	6.00	0.86	0.75	0.80
	protégé	5.00	5.00	1.00	0.83	0.91	6.00	4.00	0.67	0.50	0.57
User 4	groupie	8.00	6.00	0.75	1.00	0.86	8.00	7.00	0.88	0.88	0.88
	protégé	5.00	5.00	1.00	0.83	0.91	6.00	4.00	0.67	0.50	0.57
User 5	groupie	5.00	5.00	1.00	0.83	0.91	5.00	5.00	1.00	0.63	0.77
	protégé	5.00	5.00	1.00	0.83	0.91	6.00	6.00	1.00	0.75	0.86
User 6	groupie	6.00	5.00	0.83	0.83	0.83	6.00	4.00	0.67	0.50	0.57
	protégé	4.00	4.00	1.00	0.67	0.80	3.00	1.00	0.33	0.13	0.18
User 7	groupie	8.00	6.00	0.75	1.00	0.86	8.00	6.00	0.75	0.75	0.75
	protégé	8.00	6.00	0.75	1.00	0.86	10.00	8.00	0.80	1.00	0.89
User 8	groupie	5.00	5.00	1.00	0.83	0.91	6.00	6.00	1.00	0.75	0.86
	protégé	5.00	5.00	1.00	0.83	0.91	6.00	6.00	1.00	0.75	0.86

ROMEO AND JULIET

	condition	# entities	correct entities	entity precision	entity recall	entity f1	# properties	correct properties	property precision	property recall	property f1
	Gold	14.00					8.00				
User 1	groupie	9.00	9.00	1.00	0.64	0.78	4.00	4.00	1.00	0.50	0.67
	protégé	6.00	6.00	1.00	0.43	0.60	3.00	3.00	1.00	0.38	0.55
User 2	groupie	14.00	5.00	0.36	0.36	0.36	12.00	0.00	0.00	0.00 -	
	protégé	14.00	8.00	0.57	0.57	0.57	13.00	2.00	0.15	0.25	0.19
User 3	groupie	7.00	7.00	1.00	0.50	0.67	4.00	4.00	1.00	0.50	0.67
	protégé	8.00	8.00	1.00	0.57	0.73	4.00	2.00	0.50	0.25	0.33
User 4	groupie	8.00	8.00	1.00	0.57	0.73	6.00	6.00	1.00	0.75	0.86
	protégé	5.00	5.00	1.00	0.36	0.53	2.00	2.00	1.00	0.25	0.40
User 5	groupie	7.00	7.00	1.00	0.50	0.67	3.00	3.00	1.00	0.38	0.55
	protégé	9.00	8.00	0.89	0.57	0.70	3.00	2.00	0.67	0.25	0.36
User 6	groupie	11.00	7.00	0.64	0.50	0.56	5.00	2.00	0.40	0.25	0.31
	protégé	6.00	2.00	0.33	0.14	0.20	3.00	0.00	0.00	0.00 -	
User 7	groupie	11.00	11.00	1.00	0.79	0.88	6.00	6.00	1.00	0.75	0.86
	protégé	11.00	11.00	1.00	0.79	0.88	6.00	6.00	1.00	0.75	0.86
User 8	groupie	8.00	8.00	1.00	0.57	0.73	4.00	4.00	1.00	0.50	0.67
	protégé	9.00	7.00	0.78	0.50	0.61	4.00	0.00	0.00	0.00 -	

Figure 11: Model quality scores for each participant, expressed as precision, recall and F1-measure separately for entities and properties for each of the two exercises in the study.

the Menu task. For the rest, however, Groupie demonstrated a slight but similarly statistically insignificant advantage. In particular, Groupie outperformed Protégé the greatest in recall, where results in both exercises were nearly significant by non-parametric Wilcoxon signed-rank tests ($p < 0.12$). The lack of statistical power likely arose from the inadequately large sample size, as well as potentially other factors including design problems in the exercises we used which we discuss in “Discussion”. We are currently working on a follow-up study involving more users and experimental conditions to better identify interactions.

Time to completion - Participants took an average of 360s to complete the Menu task with Groupie, versus 275s with Protege; for the Shakespeare task, however, participants took less time with Groupie than with Protégé 281s versus 338s. Controlling for inter-participant variation on time taken and experiment, our ANOVA analysis rejected the null hypothesis; but post-hoc tests revealed this effect was due to participant rather than client used, meaning some people took significantly longer to complete exercises than others, but either Groupie nor Protégé were significantly faster.

Interaction Log Analysis - Analysis of interactions captured via screen grabs revealed several common patterns and problems.

A number of usability problems thwarted individuals during the Protégé trials. For example, when getting started with Protegé, users were overwhelmed by multitude of buttons, options and the meaning of the icons used on these buttons. Since only a few of the buttons were actually used for the task, many users accidentally clicked irrelevant buttons and had to back out of resulting dialog boxes. This behavior did not occur in Groupie as the UI options are simplified and limited, so that there is effectively only one obvious choice in any situation.

A further difficulty with Protegé concerned the property setting dialog. For a property to be set correctly, users had to select both a property and it’s value then click the ‘OK’ button. However, no warning was given if this selection was not made, resulting in users becoming confused as their actions had no effect. This was compounded by the fact that after creating a new property, it was not automatically selected. It is not possible to do this in Groupie as all actions initiated with intent to create a new property will always result in a new property actually being created.

As Protegé does not provide a search or type-ahead interface, this made users more likely to create duplicate resources, where they should have chosen a resource which already existed. This is because users had to scan the list of existing resources by eye alone, which meant they sometimes missed the correct one. Groupie provides type ahead search functionality when naming a resource which meant that users always noticed the existing resource and choose to align to it. Additionally, this search functionality provides the ability to align to an existing resource on the wider semantic web, which is not possible in Protegé without manually locating the appropriate URI first.

In Groupie, a mistake sometimes made was confusion as to which way round a property and a value should go. Property should be the outer object (group) whilst the value should be inside, but some users had this the other way round harming the consistency of their model. Some users accidentally changed their view in Groupie without meaning to and were confused by the change of state, which did not look significantly different to the previous screen. Potential solutions to these issues are discussed in "Discussion".

Exit survey Responses to the exit questionnaire revealed that all but one user felt that Groupie was significantly easier to use than Protégé, and many commented that they found it more enjoyable. Users talked about having to memorize the correct sequence of actions to work with Protegé, while Groupie was felt to be more intuitive.

Criticism of Groupie primarily surrounded interface problems arising from minor prototype bugs and aesthetic criticisms; users felt the plain interface could look more professional as well have a more attractive visual design incorporating more color. One user specifically complained that the interface looked as if it was designed by a child. Our current work in improving the aesthetics of the UI will seek a better balance between trying to inspire creative interaction but also appeal to more "serious" audiences.

DISCUSSION

In this paper we have sought to determine whether the authoring, creation and editing of linked data could be made sufficiently easy, flexible, and fun that ordinary end-user citizens, many of whom may have never even heard of the Semantic Web could create and model arbitrary models of their choosing as easily as they might use a standard mind-mapping or brainstorming tool.

Our experimental results were mixed but promising. Our basic interaction prototype, Groupie, was considered more enjoyable and easier to use than Protégé, tool of significant maturity and in widespread use worldwide by linked data engineers. With evaluation metrics we devised to evaluate both the the quality of models and the ease of use of the interface, we found that end-user citizens could quickly and fairly effectively perform easy linked data modeling tasks with both tools, but also that they experienced considerable difficulty with more challenging, open-ended modeling exercises. Much of our ongoing and future work pertains to making improvements to Groupie to make such modeling exercises feasible. Although we did not observe statistical significance in difference between the quality of models made with the two interfaces, models resulting from the Groupie conditions were overall slightly more complete than those that were created in Protégé, resulting in a slight lead in recall. We are currently working on a follow-on study with an improved version of Groupie

There are several key directions for improvement and further study.

At the highest level, we are interested in the relationship between creativity and data generation. As noted Groupie's design is informed by brainstorming/mind mapping tools which, as cited above, have been shown to help enrich idea creation more so than without these tools' presence. Participants in our study commented on the fun they had with the tool. We would like to see whether this fun in data generation also translates into better creativity or clarity with data model generation that Groupie seems to afford. In other words, not only might our tool generate usable RDF; it may simultaneously inspire better data design than data sets developed without the Groupie model approach.

Among the immediate improvements to be made to Groupie is to reduce the likelihood of getting lost or disorientated when they clicked to change perspectives between objects and properties in the view. We believe that transition animations that lead users between views in a "zoomable-UI" inspired fashion could greatly reduce the likelihood of getting lost. Additionally, a 'breadcrumbs' trail showing the path taken through the data could make it easier for users to orientate themselves. A related issue is visual distinctiveness of entities shown at the same time. As these all look the same users may find it hard to quickly identify different entities. This could be resolved through the use of hash based coloring, to give all entities a different appearance, or even by more advanced use of pictures - such as allowing a picture to be associated with each entity. So that, for example, a house would look like a house.

Entity-predicate confusion could be resolved with placeholder labels for these objects, currently just 'Untitled,' should be changed to give a hint as to their function. For example, the group could be labelled 'Is like/does/has' and the label 'Something.'

The study itself also had some limitations. These were most apparent in the Shakespeare task, which was less well specified and more open to interpretation than the menu task. This resulted in a wide variety of models being produced, many of which deviated far from what was desired, and thus were marked in the evaluation as incorrect. Such deviation behavior was manifest equally in both Groupie and Protegé, indicating that this fell from shortcomings in the modeling exercise rather than in the tools themselves. For our next evaluation, we are considering extending the training and warm up time to give users more instruction as to what the purpose of the modeling is and give context to the activity. Additionally, it is speculated users may be more comfortable with a more recent cultural example than Shakespeare!

As we have seen there are many ways we can improve Groupie to help non-experts create their own linked data sets, but also interact with existing linked data to enrich their own products. It is notable however that even though there is such room for improvement, we see evidence of pleasure, interest and delight in participants at the opportunities to create what for most of us is a new kind of information resource.Perhaps early hypercard users/creators felt the same way. That participants wanted to continue to use the tool is a good sign towards our refinement process of the tool.

CONCLUSION

In this paper we have described a light weight tool to support casual creation of linkable web data inspired by the simplicity and effectiveness of Englebart's NLS and Apple's Hypercard. We have proposed this tool as a first step in developing a personal scale rather than corporate or government scale output of linkable data that can still be readily interlinked with and enriched by any other linkable (RDF) data sets.

REFERENCES

1. Viki: A visual wiki design: http://www.lostgarden.com/2006/05/viki-visual-wiki-design.html.
2. T Berners-Lee, J. Hollenbach, K. Lu, J. Presbrey, E. Pru d'ommeaux, and m.c. schraefel. Tabulator redux: Writing into the semantic web. 2007.
3. Tim Berners-Lee. Read-Write Linked Data, 2010.
4. C Bizer, T Heath, and T Berners-Lee. Linked data–the story so far. *International Journal On Semantic Web and Information Systems*, 5(3):1–22, 2009.
5. Douglas C. Engelbart. Design considerations for knowledge workshop terminals. In *Proceedings of the June 4-8, 1973, national computer conference and exposition*, AFIPS '73, pages 221–227, New York, NY, USA, 1973. ACM.
6. P Farrand, F Hussain, and E Hennessy. The efficacy of the 'mind map' study technique. *Medical Education*, 36(5):426–431, 2002.
7. Danny Goodman. *The complete Hypercard handbook*. Bantam Books, Inc., 1988.
8. Marti Hearts. Some thoughts on social tagging, 2007.
9. Hao-Wei Hsieh and Frank M. Shipman, III. Vite: a visual interface supporting the direct manipulation of structured data using two-way mappings. In *Proceedings of the 5th international conference on Intelligent user interfaces*, IUI '00, pages 141–148, New York, NY, USA, 2000. ACM.
10. David F. Huynh, David R. Karger, and Robert C. Miller. Exhibit: lightweight structured data publishing. In *Proceedings of the 16th international conference on World Wide Web*, WWW '07, pages 737–746, New York, NY, USA, 2007. ACM.
11. m.c. schraefel and David Karger. The pathetic fallacy of rdf. In *International Workshop on the Semantic Web and User Interaction (SWUI) 2006*, 2006.
12. Ray McAleese. The Knowledge Arena as an Extension to the Concept Map: Reflection in Action. *Interactive Learning Environments*, 6(3):1–22, 1998.
13. N F Noy and D L McGuinness. Ontology Development 101: A Guide to Creating Your First Ontology, Knowledge Systems Laboratory. *Stanford University CA*, 2001.
14. A M O'Donnell, Donald F Dansereau, and Richard H Hall. Knowledge Maps as Scaffolds for Cognitive Processing. *Educational Psychology Review*, 14(1):71–86, 2002.
15. Susan Schreibman, Ray Siemens, and John Unsworth. *A companion to digital humanities*, volume 2 of *Blackwell Companions to Literature and Culture*. Blackwell, 2004.
16. Ben Shneiderman. Creativity support tools: accelerating discovery and innovation. *Commun. ACM*, 50:20–32, December 2007.
17. Giovanni Tummarello, Renaud Delbru, and Eyal Oren. *Sindice.com: Weaving the open linked data*, volume 4825 of *Lecture Notes in Computer Science*, pages 552–565. Springer-Verlag, 2007.