Provenance for the People: An HCI Perspective on the W3C PROV Standard through an Online Game

Khaled Bachour[1], Richard Wetzel[1], Martin Flintham[1],
Trung Dong Huynh[2], Tom A. Rodden[1], Luc Moreau[2]

[1] The Mixed Reality Laboratory
University of Nottingham
Nottingham, United Kingdom
{firstname.lastname}@nottingham.ac.uk

[2] Electronics and Computer Science
University of Southampton
Southampton, United Kingdom
{tdh, l.moreau}@ecs.soton.ac.uk

ABSTRACT

In the information age, tools for examining the validity of data are invaluable. Provenance is one such tool, and the PROV model proposed by the World Wide Web Consortium in 2013 offers a means of expressing provenance in a machine readable format. In this paper, we examine from a user's standpoint notions of provenance, the accessibility of the PROV model, and the general attitudes towards history and the verifiability of information in modern data society. We do this through the medium of an online-game designed to explore these issues and present the findings of the study along with a discussion of some of its implications.

Author Keywords

Provenance, PROV standard, serious game, user study

ACM Classification Keywords

H.5.2. Evaluation/Methodology, User-centered design

INTRODUCTION

In the information age, data is abundant. It is constantly created, modified, combined, deleted and manipulated in every way at rates that up until a few decades ago were unimaginable. The data, disseminated over the Internet and the World Wide Web, is forming an ever-changing record of our collective history. This history is only as valid as the data that comprises it. However, with every piece of data having its own life-cycle, its own sources and its own influences, the mesh of interdependence between bits of information on the web is prohibitively complex for anyone wishing to examine the integrity of the data before them.

Provenance, a record of the history of an object or a piece of data, is key to evaluating the validity of information. Creating a machine readable format for provenance data

would allow machines, being more suited for complex tasks, to help with the issues of data interdependence on the web. The PROV Data Model proposed by the World Wide Web Consortium (W3C) in 2013 sought to create machine readable inter-operable encoding of the history of data on the web. PROV addressed the technical challenge inherent in defining machine readable provenance.

However, looking at provenance from a user perspective is equally – if not more – important. How can provenance be stored and displayed so that is easily understood and digested by humans?

This paper is an exploration of the growing development of provenance systems and how they can be used to make the history of data and information accessible to both humans and machines. One challenge lies in finding a way to engage people with provenance and enable them to focus on its background issues. Our solution is the development of an online game-with-a-purpose that is framed around historic accounts supported by provenance. By engaging with the game, players learn about the concept of provenance in general and the PROV Model in particular and interact with it as part of the game mechanics. By placing provenance in a gaming context we hoped to create an intrinsic motivation for the players thus getting an insightful view into the minds of potential users of provenance [9].

Our contribution in this paper is an exploration of the relation people have with history at a personal and community level, how this is captured and used and how people may potentially exploit the PROV standard, a representation of provenance that was essentially designed to be processed by machines.

BACKGROUND

The study presented in this paper is aimed at investigating prevalent attitudes towards history and historic records by exposing non-expert users to notions of provenance, specifically as prescribed in the PROV Data Model. In this section we will discuss the underlying concepts of provenance and related work in analyzing the use of provenance. As we have used a game to engage people with provenance, we will also discuss other examples of using games as a tool for collecting research data.

Provenance and the PROV Data Model

The W3C defines provenance as "a record that describes the people, institutions, entities, and activities involved in producing, influencing, or delivering a piece of data or a thing" [11]. As such, provenance is generic information that captures what happened, either in a computer application or in the real world. It offers the means to verify information and infer its quality, to analyze the processes that led to a thing, and to decide whether or not it could be trusted [12]. The recent emergence of provenance as an important concern in various applications (e.g. establishing accountability, reproducibility and trustworthiness of information) had led to the release of the PROV Standards in 2013 [10]. As such, PROV is a de jure standard that is gaining traction as a means to express provenance data. Indeed, the 2014 National Climate Assessment (NCA), a four yearly report on climate change published by the US Government, employed PROV to provide traceable accounts and to support reproducibility for all of its contents [13]. Similarly, in another notable example, the Gazette (https://www.thegazette.co.uk/) – the UK's official public record since 1665 – describes the capture, transformation, and publishing processes of all its Notices using PROV. In order to make their provenance accessible to the majority of the public, who are typically not familiar with PROV, the NCA represents provenance information in a textual format (e.g. "*This image was derived from dataset nca3-...*"), while the Gazette shows a graph detailing the workflow involved in the production of a Notice. However, no information on how usable those PROV representations are to their readers is available. This is one of the main aims of the study in this paper.


Figure 1. The Entity-Activity-Agent model proposed by the PROV standard. [11]

Provenance Usability

Given that provenance information is typically complex, most efforts to improve its comprehensibility have been on devising easy-to-understand ways to represent it to end-users. As the main purpose of provenance recording is to track influences to the generation of an artifact, provenance information typically contains a number of elements and their relationships. This renders provenance information naturally suitable to the graph representation, called a *provenance graph*, whose nodes represent the elements and whose edges the relationships. There are three different types of elements in PROV [11]: Entities, Activities and Agents. An *entity* "is a physical, digital, conceptual, or other kind of thing with some fixed aspects; entities may be real or imaginary." An *activity* "is something that occurs over a period of time and acts upon or with entities; it may include consuming, processing, transforming, modifying, relocating, using, or generating entities." Finally, an *agent* "is something that bears some form of responsibility for an activity taking place, for the existence of an entity, or for another agent's activity." Between the entities, activities, and agents, there can be a number of different types of relations (as shown in Figure 1). The W3C Provenance Working Group suggested that the default shapes for each type of node in a provenance graph are: entity – ellipse, activity – rectangle, agent – pentagon [23].

Predating PROV, provenance information has already been represented as graphs. The VisTrail system [20], for example, captures the pipeline involved in the production of a visualization, i.e. its provenance, and displays it as a workflow. Another example, Probe-It allows user to query a small chunk of provenance from a mapping application and to have the information represented in a direct acyclic graph [18]. The reported trial suggested that visualizing provenance as a small graph helped its users, who were scientists, identify and explain data imperfections. The study, however, targeted participants who are experts from a narrow field, not the wider public.

More broadly, graph representations and visualizations in general have also been the subject of earlier work exploring effective ways to present graph data to end users [7]. One study, for example, investigated the usability of various visual cues to facilitate understanding of directed graphs, similar to the common directed graph representation of provenance [8]. Although those earlier studies have similarities with ours, our focus is on aspects of graphs that are unique to provenance: the accessibility and readability of the PROV standard to the general public, its suitability to encode and convey provenance information, the nature of its role in establishing trust and information confidence, its real world applications and the social and ethical implications of its use in sensitive matters such for storage of personal data.

To our knowledge there has been little work with regards to the non-expert use of provenance data. An exception is a study [1] in which the history of a computer document (e.g. which application opened, saved, renamed, and deleted it) is tracked automatically. The history could then be presented to the users in a timeline highlighting actions on documents using color-coded lines and ellipses. The study found that the timeline visualization helped users find their documents and understand their work patterns. This system, however, did not record the relationships, say, between documents, as in typical provenance graphs. In a very similar application, the Leyline system [6] additionally captures contents being "cut and pasted" between applications, and thus, is able to infer relationships between

documents tracked by it. Its provenance graphs are exposed to users as a tool to design search queries for files on desktop environments. The graphs used application icons (Word, Powerpoint, Excel, etc.) to depict the documents. The focus of this work, however, lied on the usability of the Leyline system in creating effective queries, not on the provenance graphs themselves. Nevertheless, by showing that participants were successful in creating queries with provenance graphs, the work seems to imply the usability of these graphs in general. Our work delved further into this issue by exploring various facets of user interaction with provenance such as different preferences for graph layout and different attitudes towards the use of provenance data in the real world. We approach this issue through the medium of an online game as a means to expose users to the relevant concepts as well as collect data on their use of provenance graphs.

Games and Research

Recently, there has been a significant increase in the use of games and game-like activities in a research setting, either to gather data, to affect behavior or some other form of change, as a medium for deploying and studying new technologies and activities, or more fundamentally as an area worthy of study in and of itself.

One approach to using games for research is in the form of crowdsourcing or, more specifically, a form of crowdsourcing known as human computation, in which participants are asked to perform tasks that are currently too difficult for computers. These games generally present a series of discrete tasks for players to perform and result in the generation of scientific data sets pertaining to the tasks completed. Many of these games have historically been image based. For example, the ESP game [22] requires players to tag images during gameplay under the auspices of attempting to "read the other player's mind," while others explicitly present a scientific problem to be solved, for example the identification of galaxy shapes in images in Galaxy Zoo [15] or protein folding solutions in Foldit [3].

Another approach is the rise in what are generally termed serious games, or games with a purpose specifically beyond just entertainment. Generally, all of these games serve to act as a mechanism to engage players with a conventionally non-playful situation in a playful manner, for example fitness and health [19], societal awareness [17] or even for provocative political awareness [16].

Finally, the HCI community has regularly used the deployment of novel games or game like experiences, particularly performance-led research conducted 'in the wild' [2] as an approach to studying broad interactional phenomena [21].

Fundamentally these approaches must provide a suitable motivation for players to meaningfully engage with them. Citizen science applications such as Galaxy Zoo are arguably self-motivating, making use of a player's altruistic

interest in solving the problem itself, whereas other problems, perhaps seen as less inherently worthy, require gamification [4] – the addition of external game like structures such as points or leaderboards in order to sustain interest. Previous work has suggested that this external gamification ultimately provides poor ongoing motivation for completing a series of tasks [14].

Our aim is to motivate players in engaging in provenance in a way that will not only extract data, but that is also sufficiently internally motivating that it allows ongoing investment and reflection by players, in order to allow us to probe attitudes and understanding of the subject matter.

With this in mind we developed a game around the notion of history and provenance; however, rather than presenting these concepts as a series of human computation tasks, or adding external gamification, our aim was to provide a complex game scenario in which provenance was a fundamental and playful mechanic of the game and uncovering history its goal. In other words, the main narrative thread was revealed to the player through provenance graphs.

METHODOLOGY

We developed an online game about the manipulation of history utilizing provenance graphs in the format suggested by the W3C PROV Standard as the main game mechanic. In this section we describe the details of the game, the methods of data collection and the level of participation the game was met with.

The Game

The Apocalypse of MoP (AoM) is an online game about history that uses PROV standard compliant graphs as a core gameplay mechanic. It was developed in collaboration with a local group of artists who have experience in the design and execution of alternate reality games. The underlying narrative of AoM places the player in an Orwellian version of our own reality where detailed information about every aspect of peoples' lives is documented, maintained and reviewed by a super-governmental organization called the Ministry of Provenance (MoP). The player signs up to the game by joining the ranks of an underground resistance movement named Cr0n, led by an enigmatic character known as the Groundsman (Figure 2).


Figure 2. The Groundsman (left) is the leader of Cr0n and appears in video briefings at the start of every mission, and a screenshot of the PROV orientation video (right).

The players' first mission is to infiltrate MoP by joining their "Citizen Helper" program, a crowdsourcing effort by

the Ministry that allows any citizen to help maintain the integrity of their massive database of PROV records by examining these records and highlighting errors in them. Consequently, players are introduced to the notions of provenance and the PROV model through a combination of textual guides, an interactive tutorial and orientation videos in the style of a vintage public information film (Figure 2).

After the tutorial, the game proceeds in two parallel tracks: on the one hand, the players need to maintain the image of a dutiful Ministry employee by completing simple PROV based tasks and gaining access to higher clearance levels. On the other hand, players must use their position at MoP to leak documents to Cr0n and help unmask MoP's plot for global domination. Each mission the player completes unlocks a piece of an overarching narrative, either revealing more of the Ministry's dark plot to rewrite history or pushing the story forward toward a final resolution.

The narrative which spanned four episodes was released over a period of six months from the initial launch of the game. After that, while there was no new content available, players were free to continue their duties at MoP and catch up on missions they have yet to complete at Cr0n.

The game was advertized initially through a live interactive experience at a local games festival, followed by the use of flyers and posts in relevant online forums, mainly ones frequented by fans of alternate reality games.

Infrastructure
The game ran as two separate websites: the Cr0n website, where players could complete missions and discover the game narrative through multimedia content; and the MoP website which has a lackluster office intranet design and uses a tediously bureaucratic interface where nothing can be done without filling in the necessary forms. Players can communicate with in-game characters such as the Groundsman on the Cr0n site or Sandy Spencer, the Ministry's orientation officer whose role is to provide support for players. The websites ran as two front ends of a single Django application.

Both websites offered the player an interface for examining provenance records. This was implemented with HTML5/JavaScript using the KineticJS library for the interactive elements. While visually different, as will be shown later, they are functionally identical: players could drag nodes around a canvas and create their own graph layout. Furthermore the interfaces consisted of two inspection panes where the content of selected nodes could be displayed side by side for comparison (Figure 3).

Gameplay
While working for MoP, the vast majority of tasks the player had to complete used a basic premise: given a provenance record in the form of a PROV graph, the player must find any inconsistencies in the data. An inconsistency is usually a pair of attributes in two separate nodes that,

given how the nodes are related in the graph, contradict each other. In one type of graph for example, players had to inspect the provenance of a traffic violation charge. The graph consisted of four nodes. Two of them were Agent nodes: a member of the public and a police officer; one was an Entity node describing the charge filed against the offender and the last one depicting the offence as an Activity node. In the example seen in Figure 3, the players had to notice that the registration number in the Offence node did not match the registration number in the Charge node. The interface then allowed the players to mark these two attributes and submit a report for approval.
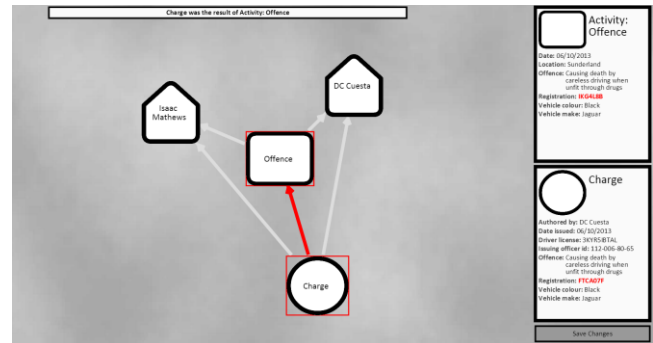

Figure 3. A basic PROV graph showing an inconsistency.

We chose this "spot the difference" style game mechanic for being simple enough while still requiring players to understand the graph in order to complete the task correctly.

New MoP tasks were created regularly and automatically by the game server based on preset templates of varying difficulty and content. The tasks were then made available to the players who could work on them after filling in the appropriate request form. Most inconsistencies required direct comparison of attribute values, such as the car registration number above; however, some were less direct in that they required some form of conversion such as from one currency to another or from an airport name to an airport code. Finally, some graphs required certain additional media to be examined, for example, looking at a luggage x-ray scan and comparing it to a written report produced by airport security personnel. To make things more challenging, every once in a while, the player would encounter a graph with no inconsistencies. During the course of the game 16 different templates were created and instantiated to provide a constant and diverse flow of tasks.

At the start of the game, players only dealt with simple graphs about everyday things such as speeding tickets and littering fines. Correctly completing tasks, rewards players with "Trust" points that they can use to request additional tasks. As players complete more and more tasks and their Trust rises, their security clearance at MoP rises as well, unlocking more complex graphs that deal with more sensitive topics such as weapons trading and orphan brainwashing programs, hinting at the dark and sinister nature of MoP's covert activities that contrasts with the image they had been maintaining.

In addition to fulfilling their role of the dutiful employees of the Ministry, players receive missions from the Cr0n website, requiring them to use the Trust they had earned to gain access to and leak documents related to Cr0n's investigation. These form the major narrative thread of the game. While most documents are also PROV graphs where inconsistencies need to be discovered, several are more generic puzzle or code-breaking tasks. Unlike MoP graphs, the provenance graphs relevant for Cr0n were created manually and pushed the narrative forward. These graphs were usually more complex than their MoP counterparts and regularly contained specifically produced media files like video or audio recordings attached to them. Spread out over 18 total missions, players completing the game would encounter 19 provenance graphs and 12 non-provenance related puzzles.


Figure 4. Two PROV interfaces used in game: the Cr0n (above) and the MoP (below) interfaces.

In keeping with the thematic difference between the organisations, MoP graphs (Figure 4) follow the basic shape model suggested by the PROV Standard on a gray background, whereas the Cr0n side offers a different feel through the use of visually distinctive elements to represent different nodes like portraits for Agents and document type icons for Entities. It also uses a cork board for a background to further mimic the feel of a detective investigation.

Participation
At the time of writing, 900 players had signed up to the game. We chose to allow players to participate in the game while opting out of the research; as such, in this paper we only look at 490 players who did consent to their data being used for research. Among these, the average age was 28, and 27% female of participants were female.

Data Collection
During the course of the game we logged all player interactions with the two websites. Special care was taken to track player interactions with provenance graphs.

After the conclusion of the game, we invited the 40 most active players in the game to conduct an hour-long phone interview, offering a £10 voucher as compensation for their time, in addition to sending an online questionnaire to all the players. Eight players agreed to be interviewed, and 41 submitted questionnaire responses. Although we targeted the most active players, half the people we interviewed had progressed only as far the first episode. This is due to a large number of players quitting the game very early on.

The interviews were semi-structured and were broken up into five phases. First, we explored the player's understanding of and familiarity with provenance. Then a walkthrough where the player was asked to solve a provenance graph using a think-aloud protocol while the interviewer observed the graph manipulation. This was followed by questions about the PROV model itself as well as the in-game interface. The fourth phase explored issues related to history, provenance and the PROV model in a real-world context. Finally we asked some questions about the game itself and the player's enjoyment of it.

The questionnaires were a combination of Likert scales and open questions exploring similar issues to the interviews. They were broken up into different sections with each section being relevant only to players who had progressed to a certain level of the game. This allowed us to get feedback from players who reached the end of the game as well as those that quit as early as the tutorial.

RESULTS
In total players completed around 13,000 provenance tasks of varying difficulty, each involving loading a PROV graph through the game interface, laying out its contents, determining what, if any, the inconsistency was and then submitting it for validation. The average time taken for a single graph was 1 minute and 46 seconds, and 96.6% of graphs were solved correctly. In the following we present the combined results from the questionnaires as well as the interviews. We have grouped the replies into different themes that deal with the idea of provenance and history in general, talk about the standard, the interpretation of the graphs and how players perceive provenance in the real world. These themes were devised from the recurring ideas in participant responses. The quotes below are derived from the interviews whereas the quantitative results are from the questionnaires. Each quote is followed by an identifier for the interviewed participant that made it.

The Notion of Provenance
The majority (68%) of respondents had not heard of provenance before participating in the game, and those that were familiar with it knew it primarily in the context of art

history and museum artifacts. Some that had particular occupations in auditing or archiving had encountered provenance in their work. Among those that were new to provenance, some distinguished knowing the term from being familiar with the concept.

"I guess it was new to me in the sense of giving it a name and defining it. Like I'd definitely thought about before what different news sources are and what their biases. But I hadn't really known a word for it." (1)

Players had no significant difficulty learning about provenance through the game. On a Likert scale, 63.4% of respondents agreed or strongly agreed with the statement "*I found learning about Provenance easy,*" and only 4.9% disagreed, with the remainder being neutral.

After playing the game, players were asked to define provenance in their own words. Most showed an understanding of the concept to varying degrees with some definitions being fairly precise:

"Provenance is keeping record of the history of something rather than just its current state, so how it's changed through time, how it's been affected by different elements along the way." (7)

However, some players' understanding of provenance ended up being very closely tied to how it is used in the game, i.e. as a tool for manipulating history by the Ministry, rather than the underlying concept, and consequently perhaps attaching to it a negative connotation:

"Provenance is about recording things and making sure it's the way they want it to be rather than how it actually is. That's what I got from it." (6)

PROV as a standard
Perhaps unsurprisingly, none of the participant reported that they knew about the PROV standard itself. They experienced it exclusively through their involvement in the game. Some found the standard fairly easy to understand:

"I thought it was really easy to understand. I had no trouble at all understanding the different terminologies involved in it." (3)

And those that didn't grasp it straight away enjoyed the challenge, seeing it as a feature of the game:

"I found it a bit hard but it was fun as well." (6)

It is important to note here that there may be a selection bias given that many players that failed to understand the PROV model even at a basic level are likely to have dropped out of the game and as such would not have contributed to the questionnaire.

Players saw the standard as "efficient", "simple", "logical" and "straight forward." When it seemed complicated, this was seen as a necessity rather than a limitation:

"I think it is complicated but it is a good way to understand complex information." (1)

Understanding Provenance Graphs
While players overall found the PROV standard to be a suitable method of modeling provenance information, there was a number of issues that were pointed out.

The most common observation offered by participants was that the direction of the arrows was seen to be counter intuitive, and in some instances confusing. This is primarily because the PROV Standard prescribes relationships that flow from one node to the nodes that influenced it e.g. *was generated by*. Visually, this meant arrows point backwards, or in chronological terms, the arrows point to the past.

"I think it's supposed to indicate that something came from something but I feel like it should go the other way. That something acted on something." (1)

"I'd say the arrows are all in the wrong order. For example, when you've got the suitcase pointing to Grover Desler, I don't think it's good English to say, 'This suitcase belongs to Grover Desler,' I would say, 'It's Grover Desler's suitcase.'" (2)

In many instances, this resulted in players completely disregarding arrowheads to avoid the confusion.

"The arrows were actually counter-intuitive or confusing to me, the direction that they went. It seemed like they went backwards from what I would expect so I never really paid much attention to the directions of the arrowheads, just the lines." (1)

This however did not limit their understanding of the graphs. Many players reported ignoring the directions of the arrows and simply treating them as non-directional links. They pointed out that the direction of the arrow often gave little to no information that could not otherwise have been deduced from the context.

"I rarely needed to look at the arrows to see what the nature of the relationship was. (...) I always saw the nodes as being linked together, rather than one pointing towards the other." (8)

Perhaps this speaks to a fundamental difference in how humans and machines process certain information: whereas machines may lack the necessary contextual knowledge to eliminate the need for explicit directionality of relationships, humans have no such difficulty.

"When I looked at the arrow from 'Scan' to 'Scanner', I didn't need the [direction of the] arrow to tell me that it was the scanner that did the scan, and when it was from 'Scanner' to 'Jonathan Cordes', I didn't need the arrow to tell me that Jonathan used the Scanner." (2)

A similar remark could be made with respect to the shapes of the nodes. Many players found the shapes unnecessary and often relied on the content of the nodes to discern the

nature of what that node represents. For example, people's names were easily recognizable and were more informative than the house-shaped node they were boxed in.

"I'm aware that there were elements, then there's people and then there's activities. But I never really when I was doing it sort them in that way. They were just nodes." (7)

"I would kind of look at what the text is in the nodes, what type of information it is labeled as. I see one looks like a name or two look like names." (1)

However, unlike with the direction of arrows, the presence of these different shapes, while not seen as particularly useful by most, did not add any confusion. Some even found them to be valuable:

"I did like the way there were different shapes to say this is a person, this is an event, this is an object, that definitely made it easier because you could focus on just the people or just the objects, and see if there was anything in those little groups that stood out." (3)

Arranging Provenance Graphs

Throughout the game, whenever a player encountered a new provenance record, the graph was collapsed such that the nodes were all piled up in the center of the screen. The player then had to manually spread the nodes out. We were interested in determining what strategy players would follow when laying out the graphs.

The W3C suggested as a convention [10] that graphs be laid out chronologically from top to bottom, left to right, such that arrows predominantly point upwards and to the left. We did not offer this recommendation to players. The resulting arrangements showed how certain people focused on different aspects of the graph. There were indeed players who opted for a chronological arrangement, believing it simplifies the understanding of the graph.

"As long as you had everything sorted in a way that made sense chronologically, it was a lot easier to understand than having a whole bunch of things at random."(3)

Others did not see a chronological arrangement as feasible.

"It doesn't work as a topological graph really, where you could do a single path." (2)

Thus, they approached the issue with different strategies, for example, focusing on grouping node types.

"First I move the nodes into groups. I put the two rectangular ones together, the two circle ones together and the three houses together." (2)

In addition to being able to foreground one particular type of node, for example Agents, when doing their analysis, sometimes an arrangement where semantically related nodes were close to each other seemed appropriate.

"I usually group like 'Scan' and 'Scanner' close to each other. 'Report' and 'Write Report.' So anything that I think

are very linked together. And people I put them in the middle, cause things usually come off people." (6)

One thing almost all players agreed on was that the best arrangements were ones where the arrows did not cross:

"First of all I'm arranging everything so I don't have any lines crossing cause it's a little bit difficult to look at everything when you have lines going every which direction." (3)

"One thing I tend to do is try and position them so that the lines connecting them don't overlap, just because I think it makes it easier to understand how everything relates." (7)

The various strategies and approaches used by players seemed to indicate that there is a rich and diverse space of potential arrangements of nodes in a PROV graph that may be suited for different types of people and different kinds of tasks that they need to do.

PROV in the real world

We asked our players to think of situations in their real lives where data in the form of the PROV Standard could be used effectively. Some of the respondents immediately saw its potential in their current or previous employment. The most salient examples being in banking to help combat fraud by supporting auditing processes:

"When you get banks being fined for colluding for example on exchange rates, then actually having an audit trail of what someone's done to change an exchange rate and who they spoke to and what actions they took and who they dealt with, and having someone check that audit trail, well that's provenance, and could save the bank millions in fines." (2)

Fraud could be combated in other domains as well:

"I keep thinking of examples of like my own work with (...) and cases of scientific fraud where if people had kept notebooks and someone who had an organized way of going through all the data that that person supposedly created or presented. It would help find where the missing or broken link was to uncover a case of fraud." (1)

Similar suggestions were made to support call centre work where managing customer relations would be a lot easier if a customer's history were described in the PROV graph structure rather than textual logs.

However, perhaps the most common use case proposed by the players was in the domain of modern day news dissemination and citizen journalism, where players referenced recent world events and how it was difficult to sort out real news from misinformation.

"Because a lot of the time I see a news article and I think, ok, I'm not sure I trust this website. (...) And then I have to spend ten minutes trying to figure out if it has a reputable source. If I had a clear idea of provenance, about who wrote it, about where they got their information from (...) that would be very useful." (8)

Personal History, Privacy and the Ethics of Provenance
Finally, we asked players about situations where they deem the collection and maintenance of provenance data to be inappropriate. It became clear that the biggest concern was linked to matters of civil liberty and personal privacy.

"Absolutely, there would be so many privacy concerns about having everything recorded and displayed. And you don't know who's looking at all this information about you when they're reviewing it. So there would be a lot of privacy concerns from this sort of thing." (3)

However, privacy concerns varied significantly with the context for which the provenance data is recorded. Workplace privacy did not seem to be a significant issue for most, with some people seeing a benefit to having these records that outweighed the loss of privacy caused.

"I wouldn't particularly mind it in my line of work because to find a mistake would be extremely helpful for me."(1)

When this type of recording did happen, it seemed important that the person recorded be aware of that.

"If my university made it clear that they were going to be checking these things, and said, here's where you sign. Here's whether you can or can't go back on it. Obviously I understand that in a situation like that there may not be the right to withdraw, but I'd want that told to me." (8)

Privacy in personal matters was treated with more caution.

"Oh good God I would not like that at all. I do some things that I'm fine with my friends and family knowing about, but I don't necessarily want the government knowing that I go to these particular websites at this time of day." (3)

Of particular concern was the nature of provenance data being a chronological record rather than individual instances of data. This brought up the issue of data longevity and the fact that provenance seems to be meant for long term if not permanent storage.

"I don't know if I would want every part of my personal life recorded forever and who would hold that information. Certainly there are ex-boyfriends I regret that I'm glad are struck from the record, said simply." (1)

Another particular feature of provenance data is that it is seen as a data format that is intrinsically less prone to error.

"For me, if the information was stored in the way you have in the graphs, I think it would probably be less likely to make mistakes." (4)

"Perhaps the only difference between sort of just interlinked data that has conclusions being drawn from and actual provenance data is how much confidence the person using the information has in that data." (7)

This is seen to emphasize both the advantages and concerns of PROV data. Privacy violations become more severe when the data is accurate and verifiable. At the same time, however, the beneficial uses of this data also become more effective.

"I would much prefer that they had enough information that those models were accurate. The one thing worse than someone building an accurate model of your life and using it to inform the way they interact with you is having an inaccurate model of you to do the same thing." (7)

Participants also indicated their preference, if PROV data about them were to be stored, on the sort of organizations they would feel most comfortable with holding this data. Governments seem to be low on the list of trusted organizations, likely due to recent developments related to some government surveillance programs.

"Well I'm not too fond of our government at this moment. I mean I would say that ideally the government would be the best place for it but we can't really trust the government to be completely impartial, so we would need to find some organization without any bias at all, which is kind of a pipe dream." (3)

The worry is that governments, like employers, may have a reason to examine one's data and take action that can be harmful for that person's wellbeing.

"I'd actually prefer that my information is kept with the Googles and Facebooks of the world. Organizations that, apart from data they have on me, have no direct effect on my life. You know Google is never going to fire me. Facebook is never going to evict me." (7)

This was not the situation for all participants, and some simply did not like the idea of giving up any amount of control over this data.

"I'd just keep it on my home computer. I wouldn't want it anywhere near the internet or anything like that." (6)

Finally, some participants raised the issue of the linked nature of provenance data, pointing out that linked data is potentially more dangerous if misused than isolated instances of data.

"The game kind of felt to me as a way of showing from the other end, how different isolated data that are quite innocuous in themselves, can be used to make very very strong and important evaluations about people and actions and situations." (7)

DISCUSSION
In this section we examine the findings presented above and discuss their implications in regards to the usability of PROV and general perceptions of history and provenance.

(1) On the suitability of the PROV model as a medium of information dissemination for the general public. While we made no attempt to examine players' understanding of textual PROV data, our experience has shown that overall players had no significant difficulty understanding graphical representations of PROV. This is evidenced by

questionnaire responses, individual player interviews as well as the high success rate in correct completion of tasks. This implies that data displayed in the form of PROV graphs could be a viable method of expressing provenance information when the target audience comprises of non-expert members of the general public. This includes graphs that contain fairly complex information within its nodes, whether it is simple text or video or audio media.

One limitation to this claim may be related to the size of the graphs. Graphs encountered by players in the game were of limited size and complexity. With the PROV standard capable of modeling extremely complex systems, it is unlikely that the same human readability would persist for very large graphs.

Players also suggested that the PROV model may be a suitable alternative to data representations currently in use in existing professional domains such as banking, customer service management and scientific research. This belief stems from both the accessibility of the model as well as its ability to express diverse and complex information as may be needed in these domains.

(2) On the directionality of arrows and the expectation of forward flow. One key issue that at times seemed to be a source of confusion for players was the backward flowing nature of PROV graphs. The model, being designed to convey information about the history of an object, suggested that arrows between nodes went from an object towards its sources of influence, rather than the other way around. Thus following arrows on PROV graphs corresponded to moving backwards in the chronology of the object. This was seen as counter-intuitive by players who had an expectation that arrows would go in the same direction as the timeline they are describing. This expectation was probably based on previous experiences with flowcharts which are typically directed forward in time. While the W3C Standard for PROV does not specify a visual representation, it does propose this directionality as a convention. Our study suggests that when the target user-base is the general public, system designers should consider reversing arrow directionality to improve the intuitiveness of the model.

That said, we do not believe that this issue negatively impacted performance because the nature of the task given to the players seemed to require little attention to the direction of the arrowheads, and often a simple presence of a link between two nodes gave sufficient information to solve the task without the direction of that link being necessary. This does not preclude there to be an issue with the readability of PROV graphs in other contexts.

Another feature of the PROV model that did not seem entirely necessary for the comprehension of the graph was the use of specific shapes for different types of nodes; however, unlike the directions of the arrows, the use of shapes was not a source of confusion. Nonetheless, this opens up the possibility of replacing these shapes with elements that are more visually appealing without compromising the readability of the graph. Icons and photos of people such as those used in the Cr0n mission graphs would create a more enjoyable experience and still maintain the expressive power of the standard. In addition, this would make it possible to identify the type of content a node describes at first glance.

(3) On the social and ethical implications of the pervasive and permanent storage of provenance. A fair deal of the player feedback focused on issues of personal privacy. Three particular things about provenance data caused serious concern when it came to privacy. First, being by definition data about the past as much as if not more than about the present, provenance data brought out some misgivings about its potential to infringe on a person's "right to be forgotten" and the right to expunge from their record certain bits of information that they no longer wish to be associated with. Second, Provenance data comes with a built-in verification mechanism and a possible context. A photo on its own might not be as much a cause for worry as that same photo attached with its complete context and a record verifying its legitimacy. Finally, seeing as how Provenance data by design links several pieces of information together in a single package it signifies another way of aggregating bits data that may have otherwise individually been harmless in isolation.

Whilst expressing concern for these issues, the participants did acknowledge that, in the right hands, provenance data can be put to beneficial uses. However, they could not deny the present potential for abuse of such a system if one were in place. Given recent developments regarding governmental surveillance programs, it is perhaps natural to wonder whether or not these abuses are inevitable.

(4) On perceptions and attitudes of history and the verifiability of historic records. While participants' attitudes on the use of provenance models to represent personal data varied from mild skepticism to deep-seated concern, its perceived potential benefits in other contexts ranged from helpful to necessary. History, in so far as it is a matter of public record, must be protected from self-serving distortion, negligent misrepresentation and everything in between. Whether it is in the context of scientific inquiry, financial trails or current events, provenance was seen as a tool that may support the goal of maintaining the integrity of historic records and thus safeguarding history in all its forms. This is in stark contrast with the attitudes towards personal histories and those related to private matters, where more and more people, with the backing of intergovernmental organizations such as the European Commission [5], insist on maintaining their right to be forgotten – the right to alter the records of their own past.

CONCLUSION

This paper explored, through the medium of an online game, some of the issues linked to people's perceptions of history, both personal and private. We did so by exposing players to notions of provenance and the PROV standard which was proposed by the W3C as a means of encoding the history of objects and things. This exposure laid the groundwork for a deep investigation of these perceptions and the attitudes associated with them.

The PROV model itself is a de jure standard for Provenance data; thus, we sought to examine its usability, its social acceptability and its perceived risks and benefits in order to establish its viability as a de facto standard for storage of historical information. Our study revealed the promising future of the standard, being simultaneously accessible and expressive. However, users' attitudes towards and perceptions of history and historical records indicated that care must be taken when applying such particular data standards to information of private or personal nature.

REFERENCES

1. Ball, R. Don't Search, Just Show Me What I Did: Visualizing Provenance of Documents and Applications. *International Journal of Human-Computer Interaction 29*, 3 (2013), 156–168.

2. Benford, S., Greenhalgh, C., Crabtree, A., ... (2013). Performance-Led Research in the Wild. ACM Trans. on Computer-Human Interaction (TOCHI), 20(3), 14.

3. Cooper, S., Treuille, A., Barbero, J., Leaver-Fay, A., Tuite, K., Khatib, F., ... & Popovic, Z. (2010, June). The challenge of designing scientific discovery games. In Proceedings of the fifth international conference on the foundations of digital games (pp. 40-47). ACM.

4. Deterding, S., Dixon, D., Khaled, R., and Nacke, L.. 2011. From game design elements to gamefulness: defining "gamification". In *Proceedings of the 15th International Academic MindTrek Conference: Envisioning Future Media Environments* (MindTrek '11). ACM, New York, NY, USA

5. The European Commission, 2012, Article 17 "Right to be forgotten and to erasure".

6. Ghorashi, S. and Jensen, C. Leyline: Provenance-Based Search Using a Graphical Sketchpad. *Proceedings of the Symposium on Human-Computer Interaction and Information Retrieval - HCIR '12*, ACM Press (2012).

7. Herman, I., Melancon, G., and Marshall, M.S. Graph visualization and navigation in information visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics 6*, 1 (2000).

8. Holten, D. and van Wijk, J.J. A User Study on Visualizing Directed Edges in Graphs. *Proceedings of the 27th international conference on Human factors in computing systems - CHI 09*, ACM Press (2009).

9. Malone, T. W. Toward a theory of intrinsically motivating instruction. *Cognitive Science 4* (1981).

10. Moreau, L. and Groth, P., eds. *PROV-Overview: An Overview of the PROV Family of Documents*. 2013.

11. Moreau, L. and Missier, P. *PROV-DM: The PROV Data Model*. 2012.

12. Moreau, L. The Foundations for Provenance on the Web. *Foundations and Trends® in Web Science 2*, 2-3 (2010), 99–241.

13. National Climate Assessment, 2014 Report. http://nca2014.globalchange.gov/; Acc. 2014-09-22.

14. Nicholson, S. (2012). A user-centered theoretical framework for meaningful gamification. Games+ Learning+ Society, 8, 1.

15. Raddick, M. Jordan, et al. "Galaxy zoo: Exploring the motivations of citizen science volunteers." Astronomy Education Review 9.1 (2010): 010103.

16. Raessens, Joost. "Reality play: Documentary computer games beyond fact and fiction." Popular Communication 4.3 (2006): 213-224.

17. Rebolledo-Mendez, G., Avramides, K., de Freitas, S., & Memarzia, K. (2009, August). Societal impact of a serious game on raising public awareness: the case of FloodSim. In Proc. of the 2009 ACM SIGGRAPH Symposium on Video Games (pp. 15-22). ACM.

18. Rio, N. Del and Pinheiro, P. Probe-It! Visualization Support for Provenance. In G. Bebis, R. Boyle, B. Parvin, et al., eds., *Advances in Visual Computing*. Springer. Berlin, Heidelberg, 2007, 732–741.

19. Schönauer, C., Pintaric, T., & Kaufmann, H. (2011, March). Full body interaction for serious games in motor rehabilitation. In Proc. of the 2nd Augmented Human International Conference (p. 4). ACM.

20. Silva, C.T., Freire, J., and Callahan, S.P. Provenance for Visualizations: Reproducibility and Beyond. *Computing in Science & Engineering 9*, 5 (2007), 82–89.

21. Tolmie, P., Benford, S., Flintham, M., ... (2012, May). Act natural: instructions, compliance and accountability in ambulatory experiences. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 1519-1528). ACM.

22. Von Ahn, L., & Dabbish, L. (2004, April). Labeling images with a computer game. In Proceedings of the SIGCHI conference on Human factors in computing systems (pp. 319-326). ACM.

23. W3C Provenance Working Group. PROV Graph Layout Conventions. 2012. Accessed 2014-09-22. http://www.w3.org/2011/prov/wiki/Diagrams