Urbanopoly: Collection and Quality Assessment of Geo-spatial Linked Data via a Human Computation Game

Irene Celino[1], Dario Cerizza[1], Simone Contessa[1], Marta Corubolo[2], Daniele Dell'Aglio[1], Emanuele Della Valle[3,1], Stefano Fumeo[1], and Federico Piccinini[1]

[1] CEFRIEL – ICT Institute of Politecnico of Milano, Milano, Italy
[2] INDACO – Politecnico di Milano, Milano, Italy
[3] Dip. di Elettronica e dell'Informazione – Politecnico di Milano, Milano, Italy

Abstract. The Urbanopoly app is a social, mobile and location-based Game with a Purpose designed around the idea of the "monopoly" board game. Urbanopoly consumes and generates geo-spatial linked and open data; through a Human Computation mechanisms, those data are validated and enriched; new information is published as linked data. Players are thus involved in the geo-spatial data collection according to the "citizens as sensors" approach and the Urbanopoly game is the entertaining expedient to assess and improve geo-spatial linked data quality.

1 Motivation

Geo-spatial data about urban environments is more and more often being openly shared and reused in social and location-based applications. As for other open datasets, *data quality* is key: outdated, incomplete or incorrect information can seriously hamper the *fitness-for-use* of the data [1] and the usefulness of the application. However, curated datasets are expensive, because they often require manual work to collect or interlink data.

Urbanopoly – the application we describe in this paper – is a mobile, location-based and social game aimed at collecting and validating the quality of geo-spatial data according to a Human Computation approach. Urbanopoly is currently available for Android devices at `http://bit.ly/urbanopoly`.

Human Computation [2] is the paradigm to leverage human capabilities to solve tasks that computers are not yet able to properly undertake. Games with a Purpose (GWAP) [3] are employed to provide an entertaining incentive to the task solution [4]. While traditional Human Computation approaches exploit users' background knowledge or domain expertise, with Urbanopoly we argue that the direct experience and "human sensing" can play an important role in solving tasks related to the physical space, like Volunteered Geographic Information (VGI [5]) tasks.

Specifically, given an initial set of geo-spatial data, the main Human Computation objective of the Urbanopoly game is:

- To *verify* the data included in the initial dataset;
- In case of imprecise or inaccurate data, to *correct* the pre-existing data;
- To *collect* new data to complement the initial data.

2 Urbanopoly Storyboard and Gameplay

Figure 1 depicts some screenshots of the application. A walk-though video is also available at `http://bit.ly/u-video`. As in the monopoly board game, the player is a *landlord* whose aim is to create a rich portfolio of venues; those *venues* are real places in the surrounding of the player, like shops, restaurants, monuments, etc.; the venues are displayed on a map on the basis of player's current location as detected by the mobile device.

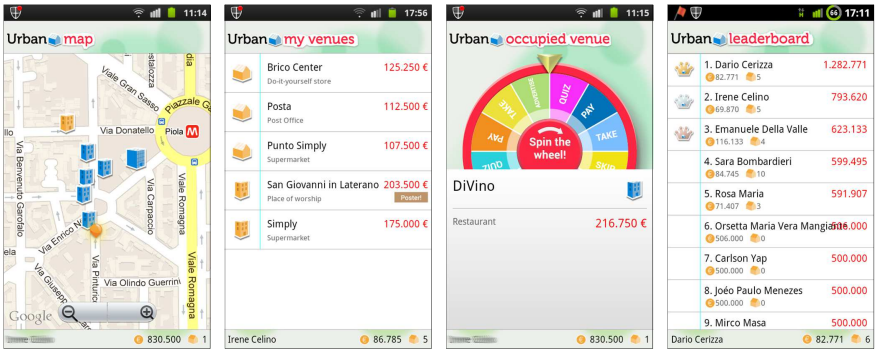

Fig. 1. Urbanopoly game storyboard. From left to right: the map showing the close-by venues, including an orange one belonging to the player; the player's venue portfolio; the "wheel of fortune" when visiting an occupied venue; the leaderboard.

Each user has an initial *budget* with which he can buy free venues. He visits a venue by clicking on a building on the map. If the venue is free and the player has enough money, he can buy it; if the venue is owned by another player, he has to spin a "*wheel of fortune*": depending on the wheel outcome, the player will be asked to face some challenges (including the hidden Human Computation tasks), which will provide him with more money. To make the game more enjoyable, however, the wheel outcome can also provide pure gaming characteristics.

As incentive to stimulate a long-term contribution, for each venue in his portfolio the player gets a "*daily bonus*" – proportional to the venue value – if he logs-in the game at least once a day. Thus, gaining money for the daily bonus and for the solution of the wheel challenges, the player can aim to the top of the *leaderboard* (cf. right-most screenshot in Figure 1). To make the game more entertaining, the plot is also enriched by other pure *gaming features*: venue trading, a mortgaging option to have immediate cash from a fictitious bank, unexpected wins or losses of money.

The monopoly-based storyboard is therefore presented to the user as a consistent game mechanics, in which the Human Computation tasks are "hidden". By pursuing the game internal objective (becoming a great landlord that owns a large number of venues), the player contributes to the achievement of the game purpose, as explained in the following.

3 Design and Development of Urbanopoly

The process followed by Urbanopoly to collect and assess geo-spatial data is illustrated in Figure 2. Starting from an initial dataset of geo-spatial information (step 1), those data are given to Urbanopoly players; by facing the game challenges, players complete some "missions" that constitute the data collection or quality assessment tasks (step 2). The "evidences" collected from players are fed into an aggregation algorithm that, as in all GWAPs, processes that input to consolidate reliable data (step 3). The verified, corrected or newly contributed information is finally published as linked open data, thus contributing back to the original sources (step 4). In the following, we provide details on those steps.

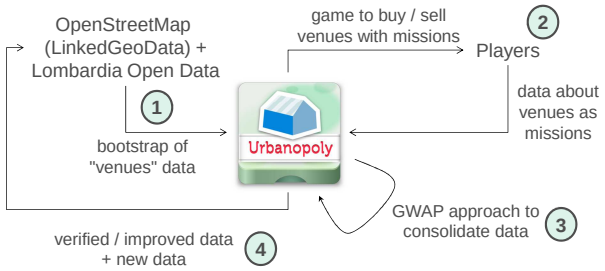

Fig. 2. Urbanopoly process to collect, verify and consolidate data.

3.1 Initial Datasets

To bootstrap the game, we took the initial information about the venues from available open data sources: a well-known VGI collaborative wiki effort – OpenStreetMap[4] – and, for what regards the area around the city of Milano, the Open Data Portal[5] of "Regione Lombardia", the local regional public authority. In the initial dataset, geographic coordinates of venues are considered stable and are employed in the map display view; all other properties of venues are collected or validated through the game.

Data from OpenStreetMap were obtained through LinkedGeoData [6], the linked data version of this VGI dataset. We selected a subset of classes representing points of interest (POIs) like shops, monuments, public transportation stops, etc. For each class we selected a number of properties to describe venues' features that Urbanopoly players can provide in the game: apart from name, category and basic address information – which are common features for all venues – restaurants are described by the cuisine type, bus stops by the line numbers, banks by the availability of an ATM, etc. Similarly, we retrieved from Lombardia Open Data Portal the information about "agriturismo" venues, quite popular in

[4] Cf. `http://www.openstreetmap.org/`.
[5] Cf. `https://dati.lombardia.it/`.

Italy. For each venue of this type, we collected their properties, like the services and products they offer to tourists.

This bunch of information is thus ready to be used in the game: the aim is to gather a high-quality set of triples in the form `<venue> <feature> <value>` (indicated as VFV in the following) in which the feature is the property and the value is its filler w.r.t. the venue. When the value is missing in the initial dataset, a Human Computation "collection" task is generated; on the contrary, when the value is already present in the dataset, a Human Computation "validation" task is proposed to the player.

We initially included 36,897 venues from Lombardia in Italy, then we added 6,749 venues from the Amsterdam area; finally, for the Semantic Web Challenge, we included also 7,817 venues from Boston for a total of more than 50,000 venues.

3.2 Human Computation Tasks

Different assignments represent different challenges for the users involved in their solution [7]; therefore, Urbanopoly was designed as a multifaceted game, in which each Human Computation task is encapsulated in a *different "mini-game"*.

Moreover, in designing those mini-games, we took also into consideration the *context* in which the users play: via a mobile phone and potentially in mobility, i.e. while moving in the urban environment. For this reason, we designed the mini-games to be easily played on a small touch screen and minimizing the requests for typing a long text (e.g. by providing – whenever available – a list of pre-defined options from which to pick the desired value).

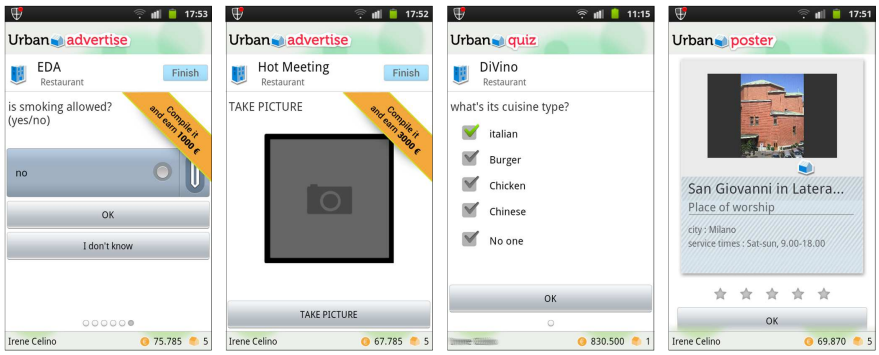

Fig. 3. Mini-games in Urbanopoly for task solving. From left to right: data collection challenges as contributions to an advertising campaign (inserting a value and taking a picture); data validation challenges to check pre-existing data or other players' contribution (answering a quiz or rating a poster).

The Human Computation tasks and the respective mini-games within the Urbanopoly gameplay are as follows (cf. also Figure 3):

– *Data collection* tasks are encapsulated in *creative* challenges: the player is asked to provide a number of elementary inputs to create a complex artifact;

- *Data validation* tasks are presented in two forms:
  - *Quiz* challenges: the option to be validated is presented together with possible alternative values and the player is asked to select the "right" answer;
  - *Rating* questions: a set of information about the same entity is presented as a single artifact and the player is asked to judge that artifact on a rating scale;
- *Data ranking* tasks are again proposed as *quizzes*: some alternative options are presented to the player who is asked to select the "best" one (e.g. the most representative depiction of a place between two or more photos).

3.3 Data Consolidation

The evidences collected from players in the game are then consolidated with the aggregation techniques [8] typical of Human Computation approaches. This consolidation lets the reliable data surface and the incorrect data be discarded.

The aggregation algorithm we designed is a weighted majority vote: to trigger the algorithm, there need to be at least two different players providing their contribution about the same VFV triple. Each player contribution is then weighted by three factors: *(1)* trouble/arduousness to collect the contribution (e.g., typing vs. check box selection, difficulty to procure the piece of data); *(2)* player's reputation (e.g. number of errors); *(3)* player's distance to the venue at contribution time (as sensed by the mobile device).

Urbanopoly uses the weighted contributions, as computed by the aggregation algorithm, to alter a *confidence score* that annotates each VFV triple. When this score becomes greater than an upper threshold, the VFV is consolidated as a correct information element; in the same way, when the confidence score becomes smaller than a lower threshold, the VFV is again consolidated, but in this case as an incorrect piece of information.

Moreover, the discovery of "wrong" information is used to penalize cheating users (with contingencies in the game and with a reduced reputation in the aggregation algorithm). To discourage malicious behaviours, an "I don't know" option is always provided when asking questions.

3.4 Publishing of Consolidated Data

The consolidated information are then given back to the community as linked data[6]. We decided not only to publish the "correct" information – by asserting the respective VFV triple – but also the "wrong" information as well as some details about the VFV confidence and the consolidation process (e.g. the contributing players, the consolidation time, etc.).

To this end, we modelled a Human Computation ontology[7] to represent those data. This ontology specializes the Provenance ontology (PROV-O [9]), currently under standardization at the W3C. This choice was made because the details about the aggregation and consolidation process describe the "provenance" of the

[6] Cf. `http://swa.cefriel.it/linkeddata/`.
[7] Cf. `http://swa.cefriel.it/ontologies/hc`.

consolidated VFV triples. In our ontology, players are the *agents*, the Human Computation tasks and the executions of the aggregation algorithm are the *activities*, and the players contributions as well as the consolidated data are the *entities* used or generated in the Human Computation process.

4 Evaluation and Outlook on Urbanopoly

We evaluated Urbanopoly using several metrics to assess different aspects of the game. With regards to the *precision* of the consolidated data, we manually checked the correctness of the output of the aggregation algorithm, one month after the release of the game on Google Play. The computed accuracy is around 92%, which appears to be a very good result.

To measure the *engagement* potential of the game, we considered a typical Game with a Purpose metrics, the Average Life Play (ALP [3], computed as ratio between the total played time and the number of active users). Urbanopoly ALP is around 100 minutes, which means that players enjoyed the game very much and returned several times to play the game. On the contrary, the game throughput in terms of Human Computation tasks solved per time unit can be improved from the current 5 consolidated VFV statements per hour; however, this can be motivated by the high number of venues and the scattered contributions of players, since the throughput in terms of collected information is of 287 evidences per hour (1,115 venues were visited at least once, but only 65 venues have at least one consolidated VFV triple).

We measured also Urbanopoly "*sociability*": playing to the game requires to log-in with a Facebook account and game achievements can be published and visualized on Facebook pages as "stories" and "story impressions"; thus, we also rely on Insight, the analytics platform of the social network. Figure 4 illustrates a statistics example, showing the Urbanopoly spreading on Facebook.

Fig. 4. Statistics on Urbanopoly achievement sharing from Facebook Insight.

Finally, to assess the "*playability*" of the game, we set up an evaluation survey at `http://bit.ly/u-survey`. Questions regard a number of different characteristics like usability, social aspects, physical presence, motivation, etc. The first feedbacks we received are very encouraging, both in terms of the general appreciation of the game and relatively to the geo-spatial dimension: players confirmed that the proximity to the venues facilitates the data collection and that they even discovered interesting facts about their surrounding environment.

Acknowledgements

This work was partially supported by the PlanetData research project, co-funded by the European Commission (FP7-257641).

References

1. Knight, S.A., Burn, J.: Developing a framework for assessing information quality on the World Wide Web. Informing Science **8** (2005)
2. von Ahn, L.: Human Computation. In: ICDE, IEEE (2008) 1–2
3. von Ahn, L.: Games with a Purpose. IEEE Computer **39**(6) (2006) 92–94
4. Guy, I., Perer, A., Daniel, T., Greenshpan, O., Turbahn, I.: Guess Who? Enriching the Social Graph through a Crowdsourcing game. In: CHI. (2011) 1373–1382
5. Goodchild, M.: Citizens as sensors: the world of volunteered geography. GeoJournal **69** (2007) 211–221
6. Stadler, C., Lehmann, J., Höffner, K., Auer, S.: LinkedGeoData: A Core for a Web of Spatial Open Data. Semantic Web Journal (2011)
7. Doan, A., Ramakrishnan, R., Halevy, A.: Crowdsourcing systems on the World-Wide Web. Commun. ACM **54**(4) (2011) 86–96
8. Law, E., von Ahn, L.: Human Computation. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers (2011)
9. Belhajjame, K., Cheney, J., Corsar, D., Garijo, D., Soiland-Reyes, S., Zednik, S., Zhao, J.: PROV-O: The PROV Ontology. W3C Working Draft (2012)

Appendix – Addressing Evaluation Requirements

Table 1 and Table 2 summarize how Urbanopoly addresses the minimal and the additional requirements (L = low, M = medium, H = high), as they are listed in the Semantic Web Challenge Criteria.

Criteria	Rating	Motivation
End-user application	H	Urbanopoly is intended as a mobile application for any citizen wishing to play with the surrounding environment
Information sources		
- diverse ownership/control	H	data sources are under different ownerships (OpenStreetMap, Lombardia Open Data)
- heterogeneous	M	data are syntactically and structurally heterogeneous (POI descriptions from different datasets)
- real-world data	H	all data used in Urbanopoly are real and considered at real scale
Meaning of data		
- Semantic Web techs	H	data are stored and published exploiting several Semantic Web technologies (e.g. triple stores, Linked Data publisher tools)
- data processing	H	Urbanopoly checks the correctness of existent data and computes new data describing the considered POIs
- alternative technologies	M	Semantic Web technologies are exploited to generate the collection/validation missions and to publish the consolidated results

Table 1. Minimal requirements

Criteria	Rating	Motivation
attractive and functional Web interface	H	the Urbanopoly interface is designed to be intuitive and easy to use, as well as pleasant and enjoyable, hiding the complexity of the technology behind it
scalable application	H	the application can scale: (1) new cities and regions can be easily added (we proved it by extending the game to Amsterdam and Boston); (2) other geo-spatial datasets can be added, both for validation and collection
rigorous evaluation	H	accuracy, engagement metrics (ALP, throughput), playability and social aspects were evaluated as reported in this paper; those metrics are continuously monitored
novelty	M/H	Semantic Web GWAPS are not new, but in Urbanopoly Human Computation tasks aim to address geo-spatial Linked Data Quality by exploiting players' mobility and physical presence in the urban space
beyond pure information retrieval	H	Urbanopoly provides data validation and data collection functionalities; the consolidated data are also published as linked data
commercial potential	M/H	Urbanopoly increases the accuracy and recency of VGI datasets; Urbanopoly can also be extended to collect other relevant data for urban stakeholders (e.g. origin-destination matrix)
ratings or rankings	H	Urbanopoly aggregates players' contributions to compute the confidence (rating) of the statements describing the POI features
use of multimedia	L/M	as part of the game challenges, players can take pictures of the surrounding venues, as well as evaluate photos taken by other players
dynamic data	H	Urbanopoly is designed around data dynamics: data collected today can become outdated tomorrow, thus Urbanopoly continuously collects and validates the geo-spatial data
results accuracy	H	our evaluation shows that the accuracy of the consolidated data is very high, over 90%
multiple languages and accessibility	L/M	the user interface is available both in English and in Italian; other languages can be easily added

Table 2. Additional Desirable Features