

# Speculation and Bluff for Linguistic Field Research – Social Software and Media for Data Gathering

Fabian Kneissl and François Bry

Ludwig-Maximilian University of Munich  
Institute for Informatics  
fabian.kneissl@ifi.lmu.de, bry@lmu.de

**Abstract:** Linguistic field research relies on gathering data on how people express themselves. Traditionally, this is done by sending scientists to interview speakers and to report on these interviews by filling forms. This is a very slow and excessively expensive process which, furthermore, is often biased. Indeed, researchers' questions and reports are influenced by their –conscious as well as unconscious– preconceptions. Crowdsourcing on the Web is a promising alternative to traditional linguistic field research. This article describes Agora, a market-like social software conceived for Web-based linguistic field research. Agora is the operating system of two distinct social media, Borsa Parole and Poker Parole, both designed to gather qualitative as well as quantitative data on Italian dialects and vernaculars. A specificity of Agora is that it can be deployed both for games and non-game social media. Arguably, Agora can be used for further applications beyond linguistic field research. Its possible application in art history is discussed.

## 1 Introduction

The Web makes crowdsourcing relatively easy and rather cheap. A core difficulty of crowdsourcing, though, is to gain a sufficiently long lasting attention of enough contributors. In spite of this difficulty, crowdsourcing has a considerable potential for the gathering of data needed in science. This article reports on a social software called “Agora”, Greek for “market”, conceived as the operating system of two social media, called “Borsa Parole” and “Poker Parole”, that have been developed as Web-based alternatives to traditional linguistic field research.

Traditional linguistic field research is done by sending scientists, typically doctoral students or other low paid researchers, who interview speakers and report on these interviews by filling forms. This is very slow, excessively expensive, and furthermore often biased. Section 2 discusses the needs of linguistic field research, especially that of Italian linguistic field research for which Agora, Borsa Parole and Poker Parole have been developed.

Agora runs a play-market where people can submit symbolic goods –like dialect phrases– together with their own assessment of characteristics of that symbolic good –where or within which social group the dialect phrase is used– and compare their own assessments with those of the community. Thus, with Agora one can “sell” a symbolic good and its

characteristics. One then receives a “payment” –in form of points which can be seen as play money or as tokens of expertise– when the community agrees. Thus, Agora is no more than the operating system for play-markets on which symbolic goods can be traded with. Section 3 describes Agora in more detail. It will be explained that markets run with Agora can be rather different what makes it meaningful to distinguish between the “market operating system” Agora and the various markets run by Agora.

Agora is first used as operating system of a market called “Borsa Parole”, Italian for “word stock exchange”. On Borsa Parole, one can deal with Italian dialect and vernacular phrases as well as their geographical and social spreads. On Borsa Parole, the better phrases and their characteristics are recognised by the user community, the more successful is a “trader”. Thus, success on Borsa Parole depends on speculating on others’ assessments. Borsa Parole is described in Section 4. This section argues that speculation, precisely because of its commonly criticised aspects, is especially convenient for gathering the data striven for in linguistics field research.

Agora is also used as an operating system of a bluff game called “Poker Parole”, Italian for “word poker”. For being successful at Poker Parole, one has to submit an Italian dialect or vernacular phrase that, except for the members of the regional and/or social group commonly using this phrase, is rarely recognised by the wider community of Italian speakers. The more seldom a phrase and its characteristics are recognised, the more successful the user is who suggested this phrase. The bluff consists in knowing a phrase’s characteristics that many other users cannot recognise and thus challenge the knowledge of other users. Poker Parole is described in Section 5. This section argues that Poker Parole unveils precisely the kind of linguistic data that Borsa Parole does not gather much. Thus, even though they both are run with the same software, Agora, Borsa Parole and Poker Parole are distinct social media.

Borsa Parole and Poker Parole fulfil distinct and complementary goals. The one social medium gathers dialect and vernacular phrases with widely acknowledged linguistic traits. The other social medium gathers phrases with linguistic traits that are rarely recognised outside their geographical and/or social spreads. Furthermore, the one social medium detects among the traders/players dialect and vernacular generalists, the other dialect and vernacular specialists. These specificities of Borsa Parole and Poker Parole and their complementarity are discussed in Sections 4 and 5.

Both social media Borsa Parole and Poker Parole are still under development. Their conception began in October 2010 as an interdisciplinary endeavour involving Computer Scientists, Romance Linguists and Computer Linguists. So far, this research has focused at fulfilling the need of Italian linguistic field research. This article reports on this research at the beginning of its experimental evaluation which has begun in March 2012 with the platform “MetropolItalia” going online in a non-advertised beta version –available at <http://beta.metropolitalia.org>. It is intended to publicise MetropolItalia in coming July and thus start its deployment.

Running (play-)markets, Agora gathers quantitative data. Indeed, prices and investments can be interpreted quantitatively. This distinguishes Agora from most social media used in Human Computation, in particular from the Games with A Purpose [vAD04], short

GWAPs, that gather symbolic data that can rarely be interpreted quantitatively. This specificity of Agora is of great significance for linguistic field research. The quantitative data gathered by Agora and why this is important for linguistic field research is discussed in Section 6.

A social medium can only be successful if it provides incentives for a wide audience to use it. Section 7 is devoted to discussing why Borsa Parole and Poker Parole are likely to provide with sufficient incentives for acquiring and keeping the audience needed for the data gathering these social media have been conceived for. An interesting aspect of this discussion is that Borsa Parole might, but does not need to, be perceived as a game while Poker Parole, being based on bluff, is unquestionably a game. Whether a social medium can be perceived as a game or not is important for its success. Indeed, playing games is an incentive in itself whereas participating in non-game social media is, in general, motivated by different incentives. Incentives of non-game social media can, but do not have to, be provided through gamification, that is, the addition of game-like rewards like badges or points.

Section 8 argues that Agora can be used for other applications than linguistic field research. A concrete possible application in art history is discussed in more detail so as to sustain the thesis of Agora's potential use outside linguistic field research.

Section 9 discusses related work and a short conclusion is stated in Section 10.

In essence, this article addresses the need for a new form of linguistic field research. Agora accounts for this need by providing the exploitation systems for two social media for gathering quantitative data and different kinds of data which complement each other – as with the two complementary social media Borsa Parole and Poker Parole.

The contributions of this article are as follows:

- Presentation of the market-like social software Agora capable of gathering quantitative data.
- Presentation of two social media, Borsa Parole and Poker Parole, that –though distinct in essential aspects– are both run by Agora. Borsa Parole and Poker Parole both aim at gathering linguistic data.
- A discussion of using Agora for gathering data in art history.

## **2 Crowdsourcing for Linguistic Field Research**

Linguistic field research aims at collecting language data together with multi-dimensional data of speakers of that language from areas where that language is spoken. The speakers' data usually consist in the speakers' geographic area, their social and education background, the situation in which the speech takes place –like formal or informal–, the time at which it takes place, and also whether the language used is written or spoken [Kre10].

Traditionally, such data is gathered by scientists, typically doctoral students or other low paid researchers, interviewing speakers and reporting on these interviews by filling forms. This is expensive and time consuming. Furthermore, this is often biased. Indeed, researchers' conscious or unconscious preconceptions might affect how people are selected and how their answers are written down by the linguists conducting the field research [Dav95, Laz95]. Because traditional linguistic field research is very expensive and very time consuming, mostly small scale studies focusing on rather small areas are conducted.

Through crowdsourcing many more people can be reached than through traditional field research. And as social media are accessed online, the area where a speaker resides can be anywhere. All speakers with an Internet connection can participate in Web-based linguistic field research. Biases induced by the Web platform or the speakers on the platform might occur, though. The analysis of gathered data about possible biases should be investigated in future (linguistic) research.

The Italian language is especially interesting for linguistic field research, making *Borsa Parole* and *Poker Parole* excellent means for investigating how to perform linguistic field research via social media. Indeed, the Italian language spoken today everywhere, in cities and countryside alike, and within all social groups is currently undergoing a divergence that originates in the big cities and spreads from there [Kre10]. This makes today's Italian different from languages such as German, English, or French.

During the restructuring and standardisation process which the Italian language experienced only in the late 19th century, that is more recently than most other European languages, a common language emerged out of several rather disparate dialects. However, instead of being perceived as languages for less educated people, the Italian vernaculars –that is, unstandardised language varieties– and dialects have remained in today's spoken and written language across all social groups [Kre11]. A witness of the strength of the Italian dialects is their presence on Wikipedia: There are small but lively versions of Wikipedia in about a dozen of Italian dialects. Currently, the vernaculars spoken in large Italian cities evolve. Especially, new vernaculars emerge, disconnecting metropolises from one another [Kre10].

The problems that traditional linguistic field research encounters are especially salient with the Italian vernaculars and dialects. The manifold vernaculars differ from (standardised) dialects and from each other in vocabulary, grammar, and/or pronunciation. Some distinctive features in language use are well known in whole, or major parts of, Italy like for example the use of “bon di” for greeting in some valleys in South Tyrol, or the use of “delizioso” (meaning “cute”) mainly by women [Kre10]. Other distinctive features in language use are, in contrast, known only in limited parts of Italy. For linguistic research, the rarely recognised phrases are just as important as the well recognised phrases. So far, there is not much data available concerning Italian vernaculars and dialects. Thus, the social media described in this article are likely to gain much importance in the near future –provided they succeed in attracting a large enough audience.

### 3 Agora – A Market for Symbolic Goods

Agora is a generic software for running play-markets in which a community of users can share both symbolic goods and assessments of characteristics of these symbolic goods. Specifically, Agora is used for running two social media, Borsa Parole and Poker Parole, where Italian dialect or vernacular phrases –that is sentences or parts of sentences– are “traded with”.<sup>1</sup> In other possible applications of Agora, completely different symbolic goods could be traded with, as discussed below in Section 8. Except in that section, the symbolic goods meant in this article are phrases, that is sentences or parts of sentences, in Italian dialects or vernaculars.

Agora makes it possible for a user to:

- add her own phrases to the market,
- propose assessments for her own phrases as well as for phrases proposed by others,
- review and adapt her own assessments based on assessments from other users.

The assessments of existing phrases are completed by the user in a number of rounds –we experiment with five rounds–, each centering on one phrase. Adding one’s own phrases and creating assessments of these can be done at any time.

While entering a phrase or considering an existing one, a user assesses the characteristics of the phrase –where and by whom the phrase is used– and estimates which proportion of the users are likely to assign the same characteristics as she does. This composition is displayed in Figure 1. More precisely, the characteristics consist of the geographical area where the phrase is spoken, the gender, age, or level of education of the person using the phrase. In Borsa Parole and Poker Parole, the geographic characteristics are stressed because the regional spread of Italian dialect and vernacular phrases is considered to be easier to grasp by Italian speakers.




Figure 1: Composition of an assessment: A user assesses the characteristics of a phrase together with her estimated proportion of users agreeing on the characteristics.

Analogue to a market, the characteristics of a phrase and an assessment of which proportion of speakers are likely to agree on these characteristics can be “sold” to other users and if other users “buy it”, that is, agree, the seller gains (play) money. The closer her

<sup>1</sup>So far, the game provides written sentences but an extension with spoken sentences is foreseen. This extension does not require any change in the media logic but only additional user interfaces for collecting and rendering spoken language.

estimation is to the proportion of users assigning the same characteristic, the more money she gains. If over time the proportion diverges from her estimation, she loses some part of that money. If it converges to her estimation, she gains some money. By submitting an assessment, a user cannot lose money gained through other assessments.

When a user reconsiders her assessments, for each one a summary of the other users' assessments is displayed. Based on this feedback she can adjust her assessments to fit the market. Here, the market regulates itself and users are rewarded for visiting the platform again.

In order to effectively gather data with social media operated by Agora, users are encouraged to suggest phrases themselves. This is important to enliven the media run on Agora so that they can grow both in the number of phrases gathered and in the number of their users.

The market of Agora is similar to a financial market like Wall Street in that users trade symbolic goods. Agora however differs from a financial market as follows: Agora is only a play-market, that is, no real money is involved. Furthermore, the ownerships of symbolic goods, characteristics and assessments are symbolic in the sense that several users can "own" the same "good" –phrase, characteristics of a phrase or assessment of the proportion of users likely to agree on that characteristics. Also, users cannot be indebted, that is, "own" a negative amount of money.

#### **4 Borsa Parole – Speculation for Gathering Well Recognised Phrases**

When a user visits Borsa Parole for the first time, the most probable action is assessing existing phrases. Then, five rounds each proposing a different phrase have to be completed after which a summary of the user's performance is displayed. In each round, a phrase is displayed which the user can skip –gaining no money– or assign a geographical area to. In Figure 2, such a characterisation of a linguistic phrase is displayed. In the GUI, the user can assign an area on the map where she speculates that the phrase is spoken. The user is provided guidance by being able to stepwise focus on smaller areas (see Figure 3).

After assigning an area, a user can estimate the proportion of other users assigning the same area and thus creating a speculation how well the phrase is recognised. Then, the assignments of other users for this phrase are disclosed in order to keep the user's attention. Next, the user is encouraged to highlight the word(s) from which she discovered where the phrase is spoken. Finally, the user can speculate on three other characteristics namely age, gender, and level of education of the speaker of the phrase and the next round begins. Alternatively, the user can opt to submit an own phrase together with an assessment on its characterisation or –if she already has assessed some phrases– review her assessments.

In Borsa Parole's assessments, the user can choose values of 30% or higher as an estimate of the proportion of users assigning the same characteristics as oneself. We expect to have to fine tune this value during the on-going experimental evaluation of Borsa Parole.

For being successful on Borsa Parole, one has to submit phrases with characteristics that

COME GIOCARE AREA COMMENTI PROFILO ESCI 0 PUNTI Ti sei registrato come: **il\_linguista**

**metropol**  
**italia** **BETA**

**LISTA DELLE SCOMESSE**

**Turno 1 su 5**

**Punti** 0  
**Viene da:** > Sud

Sicilia

**Per una scelta più precisa  
clicca nuovamente sulla  
cartina oppure conferma.**

*Prova a darti questa materia e come  
va, va!®*

Figure 2: Borsa Parole during the choice of an area for the displayed sentence: The sentence displayed on top means “Try to take this exam, whatever may happen!” The user is in round (“Turno”) 1 of 5 and has so far achieved 0 points (“Punti”). “Viene da” means “Comes from”, “Sud” means “South” (the area currently selected and marked with blue borders), the two left buttons are called “OK” and “Back”, the request in the red box on the left hand side means “For a more precise choice click on the map again, or confirm.” The button above the map is called “I don’t know. Skip”.

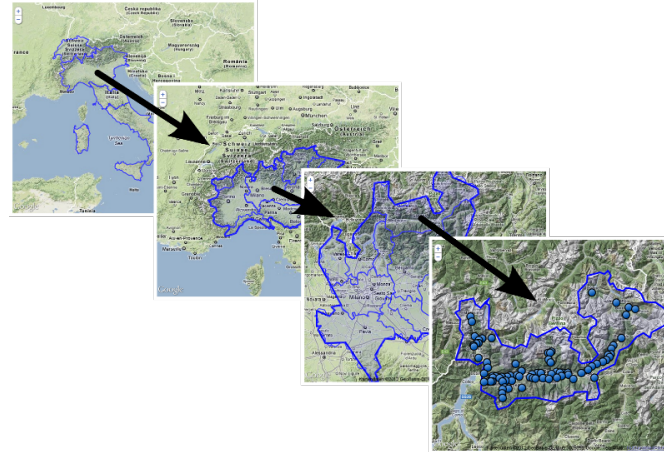


Figure 3: Stepwise focusing on smaller areas: broad geographic regions, political regions, provinces, and municipalities.

many other users of Borsa Parole are likely to agree with. As a consequence, success on Borsa Parole depends on how one is skilled at speculating on others' conceptions. This is a typical case of the "beauty contest", as Keynes described the effect in a speculative market where participants reflect on each others' behaviour and adapt their behaviour accordingly [Key36]. However, while the beauty contest analogy was meant by Keynes as a criticism of speculation on financial markets, a beauty contest-like speculation contributes to the aim of Borsa Parole. Indeed, in linguistic field research the true opinion of a single speaker is much less relevant than her perception of the community's opinion. In other words, for linguistic field research, speculating speakers are welcome!

A more detailed presentation of Borsa Parole is given in [KB12].

## 5 Poker Parole – Bluff for Gathering Rarely Recognised Phrases

Poker Parole basically works like Borsa Parole, with one exception: While success on Borsa Parole comes from submitting commonly recognised phrases, on Poker Parole it comes from submitting phrases that most users are not likely to recognise.

The gameplay of Poker Parole is the same as that of Borsa Parole but the general intention of the media is different. On Poker Parole, users are encouraged to bluff because they succeed –and thus gain money– if they submit phrases only few other people can characterise. The bluff consists in telling the community: "I guess that most of you won't be capable to correctly recognise the characteristics of the following sentence." Thus, users performing well in Poker Parole must be specialists for niche vernaculars or dialects, opposed to users performing well in Borsa Parole who must be generalists for widely known vernaculars or



dialects.

Poker Parole is tuned to favour rarely recognised phrases as follows. First, the title of the game (Italian for “words poker”) conveys the idea that users are rewarded for bluffs in the aforementioned sense. Second, the instructions displayed to the user emphasize that rarely recognised phrases are rewarded more than well recognised ones. And finally, the assessments the user can create are limited to a maximum proportion of 30% of all users assigning the same characteristic. Also here, the value of 30% might need adjustment. Therefore, the user gains most (play) money with a low level of social agreement with the other users. To prevent users from cheating by assigning a random characteristic to a phrase at least one other user must validate the characteristic of a phrase. Otherwise, no money is awarded to the creator of the assessment.

## 6 Agora Gathers Quantitative Data

Both games have in common Agora’s operating system with its data model. The type of data gathered through the games is very rich in its information value. Each assessment consists of:

- The user who trades
- The symbolic good, that is, a phrase
- The characteristics assigned to the symbolic good
- The estimate of the proportion of users who are deemed to assign the same characteristics

The possibility that several users can “own” a same symbolic good –in the sense that they create assessments for the same symbolic good– is a conscious choice. This way, meaningful results are obtained when aggregating the assessments of all users.

The proportion is a quantitative measure for the phrase’s degree of recognition. If we aggregate the proportions of all assessments for one characteristic per good, we get ratings of which goods are decisive for this characteristic, i.e., have a high average proportion, and on which goods this characteristic cannot be induced. For example, by calculating the average proportion for the characteristic of being spoken in the north of Italy for each phrase, a quantitatively stepped list of phrases is generated with the level of recognition.

Besides delivering quantitative data on the goods traded with, Agora also delivers quantitative data on the traders. The success of a user is a measure for her expertise relative to the kind of goods she trades with. If a user for example gains much (play) money with phrases from a given area, she can be considered an expert for this area.

## 7 Market or Game?

The two social media Borsa Parole and Poker Parole can only gather much and manifold data –and therefore be successful– if they provide enough incentives for users to engage on the platform. On the one hand, games are incentives in themselves. On the other hand, for non-game platforms specific incentives are needed. In the following, the incentives provided by the two social media are discussed.

First, performing well on a market is an incentive in itself. This is true for financial markets like Wall Street as well as play-markets with symbolic goods. People’s interest in performing better than the crowd is apparent in both types of markets, although in financial markets prospects of earning money play a role as incentive as well. Furthermore, each kind of market involves a gaming dimension in itself as traders are playing with each other with their trades in order to get the best performance on the market.

Concerning linguistic social media, in all cultures there is a considerable interest in language issues and in reflecting on one’s own language variations. People interested in their own language are likely to participate in Borsa Parole and Poker Parole just for seeing what others disclose on these media. No gaming dimension is needed if there are enough potential users interested in what a social medium is all about.

While the more gaming oriented Poker Parole with its bluffing is surely a game it is not so clear whether Borsa Parole with its speculations necessarily must be seen a game. There is a gaming dimension involved in that the play-market encourages speculations with phrases and play-money can be earned. But whether this is enough for being regarded as a game or merely a market is an open question. Whether the gaming aspects overweight the market aspects or not has an influence on the kinds of users engaging on the platform and also on the acceptance in the community. For example, the acceptance of Borsa Parole in the linguistic research community can lead to many researchers engaging on the market without it being a game.

Besides the conceptual game model that accounts for Poker Parole being a game also other game design patterns like game rounds or (play) money are incorporated into the two social media to lead to a more “gamified” experience. The introduction of such game design elements in non-game contexts is called gamification [DSN<sup>+</sup>11]. It has to be noted that ones refer to gamification only in non-game contexts [DDKN11]. As it is not clear whether Borsa Parole is a game or not it is unclear whether gamification might apply for it.

By allowing users to modify their speculations, the market regulates itself. A user can learn from other users and re-adjust her speculation so as to be more successful. In other words, the market incites users to express not what they might believe in but instead what they believe the community believes. This reflexivity in the behaviour of participants in a market –the beauty contest effect [Key36]– for our purposes is no deficiency but instead an advantage. Indeed, all we wish to track is how a community assesses phrases. As of language, the believes of individuals are irrelevant.

## 8 Potential of Agora in Art History

For the application area of art history, a similar application of the two complementary social media would yield new insights into the perception of artworks. The social goods traded with would be artworks and the characteristics assessed could be the artist, style, and epoch. Other than changing the GUI appropriately for displaying images of artworks instead of phrases and choosing the characteristics appropriately, the software for running such artwork-oriented social media would stay the same, that is, Agora.

In an artwork market, like *Borsa Parole*, users would speculate on the characteristics of artworks and perform well if other users assign the same characteristics. For example an image displaying the painting “Mona Lisa” would be well recognised to be created by the famous painter Leonardo da Vinci. Therefore, this market would unveil what art generalists recognise in artworks. Potential users would learn about the artists, styles, and epochs of artworks by speculating on the market.

In a complementary artwork bluff game, like *Poker Parole*, users would bet that others do not know the characteristic of an artwork. Many artworks can be mistaken for something else than what they are. For example, the authors of impressionist paintings can be confused. Such confusion are sometimes typical of artworks or painters. An artwork poker game in the style of *Poker Parole* would unveil both, such artworks and expert users of the media, yielding extremely valuable information on both, the artworks and the players, that could be used for realising an artwork search engine.

Also in other areas where there are both general and expert knowledge the complementarity of media in the style of *Borsa Parole* and *Poker Parole* are likely to be exploitable.

## 9 Related Work

A summary of what social media are and their opportunities –equally relevant for applications in research although targeted at companies– is given in [KH10]. Kinds of social media especially relevant for Agora are human computation applications [LvA11], in which humans consciously or unconsciously collaborate to solve problems computers are not able to solve. If such a problem is packaged as a game, the application is called a game with a purpose (GWAP) [vAD04]. Von Ahn and Dabbish have introduced the term GWAP for the ESP Game that solves the image labeling problem. The same image is shown to two randomly paired users who are rewarded if they suggest the same label for that image. Since the only resource shared by the two users is the image, the users tend to enter descriptions that are likely to be given also by their counterparty user. Several other GWAPs have been designed that solve different goals, among others games for protein-folding [CTB<sup>+</sup>10] and for eliciting user preferences [HvA09]. Also in art history, the GWAPs on the “Artigo” platform are employed to gather descriptive tags for artworks [SWKB11]. Suggestions for an extension of the ESP Game are given in [BW12].

Gamification is defined in [DDKN11, DSN<sup>+</sup>11] as the use of game-like design, elements, and characteristics in non-game contexts. For example, badges, leaderboards, or time con-

straints can be introduced to provide the user with motivation. Also some design decisions in Agora can be seen as gamification if the social media is a non-game – which is not clear, as discussed in Section 7.

Munro et al. present in [MBK<sup>+</sup>10] linguistic projects exploiting human computation, specifically, Amazon Mechanical Turk, where users are paid for completing tasks. An important conclusion of this article is that the linguistic quality achieved using human computation is comparable to that of controlled laboratory studies. Further articles report on using GWAPs for gathering corpora annotations [CPK08, HMS09], transcriptions [PJM07], and emotions, intentions, and attitudes [PS10]. Duolingo<sup>2</sup> is a platform offering its users support in learning languages which collects material for automating translation.

Agora, being a market, is related to prediction markets where prices are used to estimate the probability of future events [WZ04]. Agora is similar to prediction markets inasmuch as the same self-regulation is used for generating a social agreement. Agora differs from prediction markets inasmuch as, instead of predicting a future event, it serves to the collective assessment of characteristics of a symbolic good, like properties of a language by the speakers of this language. Note that some researchers have expressed the view that direct estimates might be more precise than those generated on a prediction market [Man04]. In [LBSB11] perturbations and equilibria of decision markets –similar to prediction markets with the difference that no external event is predicted but a decision assessed– are investigated.

To the best authors' knowledge, no other social media than Borsa Parole and Poker Parole have been proposed so far that rely on a market for data gathering for linguistic field research.

## 10 Conclusion

This article has described the market-like social software Agora for gathering quantitative data, two distinct social media, Borsa Parole and Poker Parole, that run on Agora have been presented and Agora's potential use for data gathering in art history has been described. An analysis whether the incentives are strong enough and all data needed are generated remains for future research.

## Acknowledgements

We thank Thomas Krefeld from the Institute for Roman Studies, Stephan Lücke from the IT Support Group for the Humanities and Christoph Wieser from the Institute for Informatics, all three with the Ludwig-Maximilian University of Munich, for useful suggestions.

This research has been funded in part by the German Foundation of Research (DFG) within the project Play4Science number 578416.

---

<sup>2</sup><http://www.duolingo.com>

## References

- [BW12] François Bry and Christoph Wieser. Squaring and Scripting the ESP Game. Research report, Institute for Informatics, Ludwig-Maximilian University of Munich, 2012.
- [CPK08] Jon Chamberlain, Massimo Poesio, and Udo Kruschwitz. Phrase Detectives: A Web-based Collaborative Annotation Game. In *Proceedings of the 4th International Conference on Semantic Systems (I-SEMANTICS)*, 2008.
- [CTB<sup>+</sup>10] Seth Cooper, Adrien Treuille, Janos Barbero, Andrew Leaver Fay, Kathleen Tuite, Firas Khatib, Alex Cho Snyder, Michael Beenen, David Salesin, David Baker, and Zoran Popović. The Challenge of Designing Scientific Discovery Games. In *Proceedings of the 5th International Conference on the Foundation of Digital Games (FDG)*, pages 40–47, 2010.
- [Dav95] Kathryn A. Davis. Qualitative Theory and Methods in Applied Linguistics Research. *TESOL Quarterly*, 29(3):427–453, 1995.
- [DDKN11] Sebastian Deterding, Dan Dixon, Rilla Khaled, and Lennart Nacke. From Game Design Elements to Gamefulness: Defining “Gamification”. In *Proceedings of the 15th International Academic MindTrek Conference: Envisioning Future Media Environments*, pages 9–15, 2011.
- [DSN<sup>+</sup>11] Sebastian Deterding, Miguel Sicart, Lennart Nacke, Kenton O’Hara, and Dan Dixon. Gamification: Using Game Design Elements in Non-Gaming Contexts. In *Proceedings of the 29th Annual CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pages 2425–2428, 2011.
- [HMS09] Barbora Hladká, Jíří Mírovský, and Pavel Schlesinger. Designing a Language Game for Collecting Coreference Annotation. In *Proceedings of the 3rd Linguistic Annotation Workshop (ACL-IJCNLP)*, pages 52–55, 2009.
- [HvA09] Severin Hacker and Luis von Ahn. Matchin: Eliciting User Preferences with an Online Game. In *Proceedings of the 27th Annual CHI Conference on Human Factors in Computing Systems*, pages 1207–1216, 2009.
- [KB12] Fabian Kneissl and François Bry. Borsa Parole - A Market for Linguistic Speculation. Research report, Institute for Informatics, Ludwig-Maximilian University of Munich, 2012.
- [Key36] John Maynard Keynes. *The General Theory of Employment, Interest, and Money*. Macmillan Cambridge University Press, 1936.
- [KH10] Andreas M. Kaplan and Michael Haenlein. Users of the World, Unite! The Challenges and Opportunities of Social Media. *Business Horizons*, 53(1):59–68, 2010.
- [Kre10] Thomas Krefeld. Italienische Varietätenlinguistik. *Italienisch. Zeitschrift für italienische Sprache und Literatur*, 63:56–62, 2010. In German.
- [Kre11] Thomas Krefeld. Alter Standard - Neue Medien. Zur Erfassung von Restandardisierungsprozessen im Italienischen. In Sarah Dessì Schmid, Jochen Hafner, and Sabine Heinemann, editors, *Koineisierung und Standardisierung in der Romania*, pages 269–281. Universitätsverlag Winter, 2011. In German.
- [Laz95] Anne Lazaraton. Qualitative Research: in Applied Linguistics: A Progress Report. *TESOL Quarterly*, 29(3):455–472, 1995.

- [LBSB11] Stephan Leutenmayr, François Bry, Tom Schiebler, and Felix Brodbeck. Work in Progress: Do They Really Mean it? Assessing Decision Market Outcomes. In *Proceedings of the 4th Workshop Digital Social Networks*, 2011.
- [LvA11] Edith Law and Luis von Ahn. Human Computation. In Ronald J. Brachman, William W. Cohen, and Thomas Dietterich, editors, *Synthesis Lectures on Artificial Intelligence and Machine Learning*, pages 1–121. Morgan & Claypool Publishers, 2011.
- [Man04] Charles F. Manski. Interpreting the Predictions of Prediction Markets. Technical report, National Bureau of Economic Research, 2004.
- [MBK<sup>+</sup>10] Robert Munro, Steven Bethard, Victor Kuperman, Vicky Tzuyin Lai, Robin Melnick, Christopher Potts, Tyler Schnoebelen, and Harry Tily. Crowdsourcing and Language Studies: The new Generation of Linguistic Data. In *Proceedings of the Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk at NAACL-HLT*, pages 122–130, 2010.
- [PJM07] Tim Paek, Yun-cheng Ju, and Christopher Meek. People Watcher: A Game for Eliciting Human-Transcribed Data for Automated Directory Assistance. In *Proceedings of the Annual Conference of the International Speech Communication Association (INTER-SPEECH)*, pages 1322–1325, 2007.
- [PS10] Lisa Pearl and Mark Steyvers. Identifying Emotions, Intentions, and Attitudes in Text Using a Game with a Purpose. In *Proceedings of NAACL-HLT Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 71–79, 2010.
- [SWKB11] Bartholomäus Steinmayr, Christoph Wieser, Fabian Kneissl, and François Bry. Karido: A GWAP for Telling Artworks Apart. In *Proceedings of the 16th International Conference on Computer Games (CGAMES)*, pages 193–200, 2011.
- [vAD04] Luis von Ahn and Laura Dabbish. Labeling Images with a Computer Game. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*, pages 319–326, 2004.
- [WZ04] Justin Wolfers and Eric Zitzewitz. Prediction Markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.