

The potential of crowdsourcing in modern lexicography

Darja Fišer and Jaka Čibej

Abstract

Due to increasing volumes of linguistic data and time constraints, the nature of lexicographic work has changed significantly in the past two decades. A number of steps in dictionary production have already been automated, but the developed algorithms are still far from perfect. Dictionary construction thus still involves a number of routine but time-consuming and expensive manual procedures for which experienced lexicographers are overqualified. This is why contemporary lexicography has started to explore options such as crowdsourcing, which can save both time and financial resources without negative effects on the quality of the results, provided that key principles of microtask design and campaign management, which are presented in this paper, are taken into account.

Keywords: crowdsourcing, microtask design, crowd motivation, quality control, legal and ethical aspects of crowdsourcing

1 INTRODUCTION

Over the past decade, the growing presence of the Internet and ever greater digitisation of work have led to numerous forms of online collaboration in which users contribute toward a common project. Aside from open-source projects (e.g. *Linux*) and collaborative initiatives (e.g. *Wikipedia*), these new forms of work also include crowdsourcing, a process in which a group of people (the crowd) contributes toward a specific goal by dividing the work load. Each individual takes on a small and manageable task that does not require much effort or time to complete, while the combined results represent a significant achievement (Howe, 2008). It is important to note that members of the crowd are typically amateurs, not field experts. Nevertheless, a number of crowdsourcing projects have shown that, with adequate support and task design, even non-experts are capable of solving tasks that were once the domain of experts. Modern technology and the wide availability of the Internet have made harnessing the potential of crowdsourcing increasingly simple and efficient.

The term *crowdsourcing* was first introduced by Jeff Howe in 2006, and has since been used to describe a wide range of work practices. In order to separate crowdsourcing from other forms of collaborative work, such as *co-creation* and *user innovation*, Estellés-Arolas and González-Ladrón-de-Guevara (2012) propose the following definition:

“Crowdsourcing is a type of participative online activity in which an individual, an institution, a non-profit organization, or company proposes to a group of individuals of varying knowledge, heterogeneity, and number, via a flexible open call, the voluntary undertaking of a task. The undertaking of the task, of variable complexity and modularity, and in which the crowd should participate bringing their work, money, knowledge and/or experience, always entails mutual benefit. The user will receive the satisfaction of a given type of need, be it economic, social recognition, self-esteem, or the development of individual skills, while the crowdsourcer will obtain and utilize to their advantage that what the user has brought to the venture, whose form will depend on the type of activity undertaken.” (ibid.: 9–10)

An essential element of a crowdsourcing project is its initiator, a company, organisation or individual who designs and manages the campaign as well as recruits crowdsourcers to perform the specified tasks. The crowdsourcers' contribution benefits both the initiator, who obtains valuable data, as well as the participants, who receive either monetary or some other type of compensation in return for their services.

In modern lexicography, the most widespread form of online user contribution is collaborative lexicography, which involves users providing new dictionary entries or suggesting updates of existing ones (Abel and Meyer 2013). The most

famous examples of collaborative dictionary projects are *Wiktionary*¹ and *Urban Dictionary*.² For Slovene, the best-known among such projects is *Razvezani jezik* (The Tongue Unleashed), but many smaller ones also exist, focusing mainly on collecting dialectal vocabularies.³

As the increasing automation of lexicographic work has been introduced to tackle stricter time constraints and increasing quantities of data, certain phases of dictionary construction have become routine tasks for which the lexicographers are overqualified. In this context, crowdsourcing has a lot of potential and can save valuable time, not as the main phase of dictionary construction, but as a way of post-processing, cleaning-up and validating automatically extracted data. It is therefore surprising that crowdsourcing has not been embraced by publishers and incorporated in the workflow of lexicographic projects. This is why the goal of this paper is to demonstrate successful implementations of crowdsourcing in related fields, as well as to outline the key principles of crowdsourcing task design and project management in lexicography.

2 CROWDSOURCING LANGUAGE DATA

In this section, we present an overview of related projects from various fields of natural language processing that have successfully implemented crowdsourcing.

2.1 Language resources

Klubička and Ljubešić (2014) used crowdsourcing to build an MSD-tagged and lemmatised corpus of Croatian to be used as a dataset. The evaluation of the crowdsourcing process showed that the accuracy of an individual crowdsourcer amounts to 90% on average, while the average accuracy of the majority vote of three crowdsourcers was approximately 97%.

Through their online application *Wordrobe*⁴, Venhuizen et al. (2013) presented crowdsourcers with a series of tasks to annotate the *Groningen Meaning Bank*.⁵ These tasks included homograph disambiguation, proper noun annotation and sense assignment to polysemous words. Compared to the gold standard, the results proved to be very reliable, even with only a small number of answers being collected.

1 <http://sl.wiktionary.org/>

2 <http://www.urbandictionary.com/>

3 <http://razvezanijezik.org/>

4 <http://wordrobe.housing.rug.nl/>

5 <http://gmb.let.rug.nl/>

Rumshisky (2011) and Rumshisky et al. (2012) used the *Amazon Mechanical Turk* crowdsourcing platform to build a semantically annotated corpus and semantic lexicon of English, both annotated by non-expert native speakers. The results showed that even non-experts can build a product with the same quality as one developed by experts.

Fossati et al. (2013) used the *CrowdFlower* crowdsourcing platform to annotate semantic roles in English texts. The comparison of crowdsourcing to standard annotation methods showed that the former, which divides the annotation process into several less complex steps, is faster as well as more accurate.

Fišer et al. (2014) used *sloW Crowd* (Tavčar et al. 2012), a custom-designed crowdsourcing tool, to clean up errors in the automatically constructed semantic lexicon *sloWNet*. With an average inter-annotator agreement of 80%, which is high for complex semantic tasks, the crowdsourcers showed a high degree of consensus, with very few ambiguous solutions.

Crowdsourcing was also used by Kosem et al. (2013) to validate and classify automatically extracted collocations and examples of use from the Slovene corpus *Gigafida*. The results of their experiment showed that a well-formulated, one-dimensional and objective question is crucial to achieving reliable crowdsourcing results.

Last but not least, the online game *Igra besed* (A Game of Words)⁶ was designed to collect collocations by asking the player to suggest three typical adjective or noun collocates for a random noun or adjective. The suggestions are then scored according to the ranked list of collocations automatically extracted from the *Gigafida* corpus. The game offers a single-player mode as well as two two-player modes, either with a selected or a random player. The game collects the players' answers as well as their metadata (e.g. usernames, time of playing, co-player, points collected). The goal of the campaign is to identify the collocability measure that best represents the speakers' language intuition.

2.2 Language technologies

Crowdsourcing in language technologies was first embraced by machine translation researchers. Zaidan and Callison-Burch (2011) recruited crowdsourcers to vote for the best machine translation suggestion, and showed that the quality of the crowdsourced results rivals the work of professional translators. Crowdsourcing has also been successfully used to evaluate machine translation systems (Bentivogli et al. 2011; Denkowski and Lavie 2010), align texts in parallel corpora

⁶ <http://www.igra-besed.si>

(Gao and Vogel 2010) and collect datasets for statistical machine translation (Negri et al. 2011).

Chamberlain et al. (2008) use the online game *Phrase Detectives*⁷ to crowdsource data for anaphora resolution. To attract more players, they have made their game available as a Facebook app.

Snow et al. (2008) conducted crowdsourcing campaigns on *Amazon Mechanical Turk* for a series of tasks, e.g. sentiment analysis in English newspaper titles. The evaluation of the data annotated by non-experts showed that as few as four annotations per task are required to achieve expert quality levels.

3 MICROTASK DESIGN

The basic concept of crowdsourcing is to divide a complex and large-scale problem into smaller, simpler and more manageable parts. The overall collection of activities that aim to provide a solution for the problem at hand is called the *crowdsourcing campaign*. The individual parts that are solved by crowdsourceers are called *microtasks*. Microtask design is a key phase in every crowdsourcing project. In this section, we provide an overview of the principles that need to be taken into account when designing microtasks in order to attain quality and useful crowdsourcing results. In addition, we provide several examples of well-designed microtasks.

3.1 Key principles of microtask design

Simplicity – Because microtasks are often undertaken by non-experts, it is important to keep the tasks as cognitively simple as possible, with the goal of collecting as many answers as possible in the shortest time period (cf. Rumshisky 2011; Snow et al. 2008; Lease and Alonso 2014).

Adequate questions – Microtasks should not include questions that cannot provide accurate and reliable results, such as opaque or ambiguous questions or overly subjective and unreliable estimates (cf. Kosem et al. 2013b). The questions posed need to be one-dimensional, which is why complex, multi-dimensional problems need to be split into several less complex steps (cf. Biemann and Nygaard 2010).

Adaptation to the target group – Different problems require varying levels of expertise. This needs to be taken into account when recruiting crowdsourceers

⁷ <http://anawiki.essex.ac.uk/phrasedetectives/>

(e.g. non-experts, students, or experts). Crowdsourcers with insufficient knowledge require more training (which is time-consuming and often destimulating) and will provide less reliable results. On the other hand, it is both demotivating and expensive to hire experts for trivial tasks.

Technical simplicity and user-friendly interface – User registration, login and task solving need to be logistically straightforward. They should not involve too much clicking or movement across the screen, and should avoid unnecessary manual data entry. This is already incorporated in most popular crowdsourcing platforms, but needs to be taken into account if using a custom-built one. Von Ahn and Dabbish (2008) stress the importance of bite-sized batches of micro-tasks that can easily be solved in a single sitting, while Jurgens and Navigli (2014) emphasise the importance of a user-friendly interface which does not rely on linguistic terminology.

Short instructions – The instructions for solving microtasks must be concise and should include illustrative examples.

Feedback – It is recommended to provide crowdsourcers with feedback for their answers. This allows them to improve and at the same time motivates them to continue their work.

Challenge, randomness and time restriction – Von Ahn and Dabbish (2008) make a number of recommendations for successful GWAP design that are relevant for other types of crowdsourcing campaigns. For example, the more the task is entertaining to solve, the more effective it is. It therefore needs to be designed in such a way that it presents a challenge to the player, e.g. by introducing scores, time restrictions, halls of fame, and so on. The task should also include an element of randomness, e.g. in assigning partners or questions. This prevents the tasks from being too predictable, while also eliminating the possibility of cheating.

3.2 Microtask examples

In this section, we present examples of various types of crowdsourcing tasks that have proved successful in related works.

3.2.1 *Classical microtasks*

Figure 1 shows an example of a microtask for semantic role annotation (Fossati et al. 2013). This consists of a brief instruction followed by a sentence in which the

crowdsourcer is asked to annotate the *agent* and *body part*. In this case, the correct answers are *he* and *none*.

Can you understand the meaning of words?

Instructions -

Please read the given sentence. It is about an event which is defined in the title and bolded in the sentence. Then read each definition and select the matching piece of text.

Warning! If you think there is **NO** matching, please answer None.

Body movement

And once he had heard Sweetheart coming down the stairs , her high-heels ringing on the stone steps , and he had **thrown** the stolen food in Rosie 's corner in a panic .

agent: the agent uses some part of his/her body to perform the action.

he
 the stolen food
 in Rosie 's corner
 None

body part: this element describes the body part that is involved in the action.

he
 the stolen food
 in Rosie 's corner
 None

Figure 1: Microtask for semantic role annotation.

Figure 2 shows an example of microtasks used for removing noise in the automatically generated sloWNet (Fišer et al. 2014b) in sloWCrowd (Tavčar et al. 2012). The crowdsourcer needs to confirm or reject an automatically assigned literal candidate (word or word phrase) to a wordnet synset using the English equivalents and definition provided. In this case, the correct answer is *no*.

Literal validation

Is the automatically translated expression "čas" a suitable lexicalization of the concept "a punctuation mark (.) placed at the end of a declarative sentence to indicate a full stop or after abbreviations"?

LITERAL
čas

SYNSET
full point, stop, full stop, point, period

DEFINITION
a punctuation mark (.) placed at the end of a declarative sentence to indicate a full stop or after abbreviations

Yes
 No
 Skip

30%

Figure 2: Microtask for synset validation.

3.2.2 *Microtasks through games with a purpose*

Figure 3 shows the interface of *Puzzle Racer* (Jurgens and Navigli 2014), a GWAP in which the player drives a race car in order to gather coins and other rewards that lead to the collection of points. Before the start of the game, the player is presented with a hint in the form of three images (Figure 4). The player then needs to find out what the images have in common in order to solve the puzzle presented during the racing phase (Figure 5). In this case, the correct answer is *money*.



Figure 3: *Puzzle Racer*, a game with a purpose.




Figure 4: Hint in the *Puzzle Racer* game.

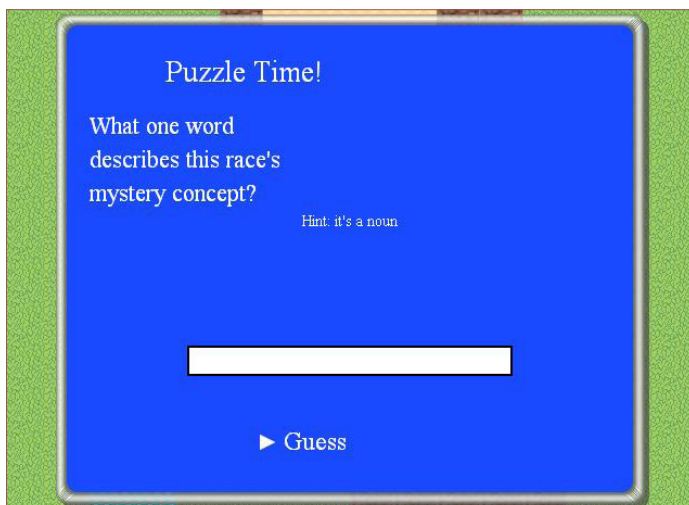


Figure 5: Riddle in the *Puzzle Racer* game.

Figure 6 shows the interface of *Igra besed* (A Game of Words). The player is presented with a word (in this case the adjective *gasilska*) and is required to provide three suggestions of typical collocations within 30 seconds. The game also has a two-player mode in which a player can compete against a chosen or random opponent. The player's answer yields points based on the word's ranking in the collocation lists from the Gigafida corpus of Slovene.




Figure 6: The interface for *Igra besed*.

3.2.3 Microtasks on social media

Games with a purpose can also be implemented on social media, where they are readily accessible to a large pool of users. Figure 7 shows a screenshot of the Facebook version of *Phrase Detectives* (Chamberlain et al. 2008). The player is presented with an example with two highlighted phrases, one of which refers to the second, and is asked to validate the correct annotations. Points are awarded according to inter-annotator agreement among the players who answered the same question. In this case, the correct answer is *agree*.



Figure 7: The Facebook version of *Phrase Detectives*.

4 MOTIVATIONAL ASPECTS OF CROWDSOURCING

Motivation is a crucial aspect of crowdsourcing projects, especially for languages with a limited pool of crowdsourcers. Crowdsourcing motivation can be either material or non-material, but it should always fulfil one or more of the crowdsourcers' needs, e.g. financial compensation, social recognition, confidence boost, or skill development. When discussing the motivation behind user-generated content on the web, Lew (2013) considers three categories of motivation: psychological, social, and economic. These are discussed in terms of crowdsourcing in this section.

4.1 Social motivation

The social aspect of motivation is based on the individual's need to connect and collaborate with other individuals sharing similar interests. This type of

collaboration enables all participants to gain new knowledge or skills and to improve their reputations in the community.

4.1.1 Community motivation

When motivating crowdsourcers, enthusiasm is much more important than the size of the community. It is important for the members to identify with the community and have the desire to contribute to its success, development, or recognition, and this has been the driving force of most collaborative lexicographic projects. For example, in the 10 years since the launch of *Razvezani jezik*, the Slovene collaborative dictionary of creative language use, approximately 1,600 anonymous authors have contributed more than 3,700 entry words and 2,300 dictionary entries (Dolar 2014). There is also substantial user involvement in language-related user groups on Facebook, which suggests that Slovenes could also be motivated to participate in the construction of Slovene language resources.

4.1.2 Educational motivation

Educational motivation is found in tasks that enable crowdsourcers to gain new knowledge or skills. This approach has been implemented by *Duolingo*⁸ (von Ahn 2013), a website offering free language courses that consist of various tasks, e.g. translating sentences into foreign languages, which are then used to train models for machine translation of web content.

4.1.3 Acknowledgments and titles

Another type of social motivation are the acknowledgments that crowdsourcers receive as a reward for their work in the community. These range from certificates to prestigious titles (e.g. Wikipedia editor), or a mention in the community's hall of fame.

4.2 Psychological motivation

Many users find creating user-generated content psychologically fulfilling, either because they enjoy sharing their knowledge with others, because this allows them

⁸ <https://www.duolingo.com/>

to fulfil a need for self-expression, or simply because they find it entertaining. One of the examples that best utilise entertainment to obtain tangible results are games with a purpose (GWAP), an increasingly popular form of collaborative work/crowdsourcing. These rely on people with an internet connection and an interest in video games to improve search engine performance, filter content, and collect linguistic data through gameplay, tasks still too complex to be performed by computers.

Two early GWAP for collecting linguistic data were the *ESP Game* (von Ahn and Dabbish 2004) and *Peekaboom* (von Ahn et al. 2006b). In the *ESP Game*, a pair of random players are presented with an image. The goal of the game is to guess the word the second player will use to describe the image. The results data has proved useful in a number of ways, such as improving search engines and developing software for the visually impaired. The second game, *Peekaboom*, employs a similar concept: the players are presented with an image and asked to determine the position of an object. The collected data is then used to train computer vision models.

Other successful games with a purpose include *JeuxDeMots* (Joubert and Lafourcade 2012), a game designed to build a French lexical network; the aforementioned *Phrase Detectives* (Chamberlain et al. 2008); *Puzzle Racer* (Jurgens and Navigli 2014), a game for semantic image annotation; and *Verbosity* (von Ahn et al. 2006a), which uses questions or sentence completion tasks to collect the general knowledge data (e.g. statements such as *milk is white*) needed to build ontologies and improve the intelligence of computer systems.

4.3 Economic motivation

Economic motivation is based on monetary remuneration or other material rewards for completing tasks.

4.3.1 Micropayments

Monetary remuneration is commonly employed in large-scale (especially commercial) projects that expect crowdsourcers to complete a significant amount of work over a longer period of time. It is usually implemented in the form of *micropayments*, paid out for each task or a batch of tasks solved. Micropayments are supported by most popular crowdsourcing platforms, such as *Amazon Mechanical Turk*,⁹ *CrowdFlower*,¹⁰ and *Clickworker*.¹¹

⁹ <https://www.mturk.com/>

¹⁰ <http://www.crowdflower.com/>

¹¹ <http://www.clickworker.com/en/>

The crowdsourcing workflow involving micropayments is structured as follows: the crowdsourcing initiator uploads their project, consisting of a set of microtasks, on a crowdsourcing platform and transfers a certain amount of funds to the platform owner. The amount depends on project size, task complexity, the number of different tasks, and so on. The platform owner is entitled to a share for hosting the project, while the rest of the amount is distributed to crowdsourcers according to the number of tasks they complete.

Micropayments have been used in a number of linguistic crowdsourcing projects (Akkaya et al. 2010; Rumshisky 2011; Rumshisky et al. 2012; Fossati et al. 2013). However, it should be noted that certain platforms, e.g. Amazon Mechanical Turk, have their own pool of crowdsourcers (registered users allowed to solve tasks) that mainly consists of English speakers (or speakers of other large languages), and therefore cannot be used by researchers interested in smaller languages. In addition, local financing and tax legislation may pose additional restrictions for project financing and micropayment management.

4.3.2 Other rewards

Economic motivation may be achieved through other types of rewards, such as vouchers, tickets, software licences, or other tangible benefits (T-shirts, pins, etc.). It is commonly employed by small-scale projects (El-Haj et al. 2014; Fišer et al. 2015) with limited funding, often in combination with social and/or psychological motivation.

5 QUALITY CONTROL AND ASSURANCE

In this section we present some of the mechanisms used for direct or indirect quality assurance of crowdsourcing results and noise elimination, in order to overcome issues related to unclear instructions or unreliable crowdsourcers.

5.1 Gold standard

The most common method of quality control is the gold standard, a manually annotated dataset of microtasks that have been solved correctly by an expert. The gold standard microtasks are randomly added to the batches of genuine microtasks in order to measure the reliability of crowdsourcers and exclude the responses of unreliable individuals from the final results.

When forming the gold standard, it is necessary to take into consideration that it should be representative of the entirety of microtasks, both in scale and difficulty. If the gold standard is too simple, it cannot effectively separate reliable and unreliable crowdsourcers. On the other hand, a too complex gold standard will exclude too many crowdsourcers. In addition, when solving microtasks an appropriate balance of gold standard and genuine microtasks is necessary. Too few gold standard microtasks cannot reliably track a crowdsourcer's accuracy, while too many will be uneconomic, as this would mean that the crowdsourcers are provided with tasks that have already been solved.

5.2 Inter-annotator agreement

The second means of assuring the quality of crowdsourced data is via inter-annotator agreement. By presenting multiple crowdsourcers with the same microtask, several answers are available for every task. The distribution of the answers can then be used to calculate a confidence score for every individual microtask or crowdsourcer (Oyama et al. 2013). In this scenario the majority vote can be taken into account when making the final decision, which means that we accept the answer provided by the majority of the crowdsourcers.

Low inter-annotator agreement might indicate that the microtasks were not properly designed, assigned to crowdsourcers with insufficient knowledge, or presented with unclear annotation guidelines. The optimal number of times the same question is given to multiple crowdsourcers is very important in this context, as each further repetition increases the costs but does not provide new answers. For most tasks three annotations are required, while more complex tasks call for 5 repetitions (Krek et al. 2013b).

5.3 Refereeing

Refereeing is a process in which problematic examples that the crowdsourcers were unable to solve reliably are resolved by an expert referee. When microtasks and annotation guidelines are well-prepared, then referees are left only with a small portion of difficult tasks to resolve, while the bulk of the work is still crowdsourced. With regard to annotating a corpus of Croatian, Klubička and Ljubešić (2014) report that this process nearly halved the amount of work done by experts.

5.4 Crowdsourcer consistency

The final method of quality control is based on the concept of consistency, also called *intra-annotator agreement* (Gut and Bayerl 2004), which tracks the consistency of a crowdsourcer's answers when presented with the same microtask multiple times throughout the annotation session. If the answers to the same task differ to a great extent, then they are excluded from the final dataset, as the crowdsourcer is either not confident or knowledgeable enough, or is selecting random answers in order to increase their productivity.

6 LEGAL, FINANCIAL AND ETHICAL ASPECTS OF CROWDSOURCING

In this section, we present the legal, financial and ethical factors to be taken into account when setting up a crowdsourcing project. The related restrictions depend on local legislation, and although they do not directly influence project quality or content, they often present an obstacle to using crowdsourcing in research, especially in lexicography, where many researchers are not familiar with legal restrictions and rarely have access to expert legal advice. Moreover, as crowdsourcing is a relatively new form of work, it is not explicitly covered by the current legislation, which is why several issues remain unresolved.

6.1 Crowdsourcer payment

Sabou et al. (2014) appeal for ethical micropayments so that crowdsourcers' earnings correspond to the local average salary or hourly wage that is standard for the services provided. Fort et al. (2014) warn that the concept of crowdsourcing as a new form of work is still absent in work-related legislation, which puts crowdsourcers in a precarious position when it comes to payment, occupational safety, worker's rights, and so on. This remains the case, even though as many as 20% of workers on Amazon Mechanical Turk are said to earn their living exclusively through crowdsourcing. It is thus imperative that crowdsourcers are guaranteed fair and prompt payment, which is not always the case in many projects (Silberman et al. 2010).

Sabou et al. (2014) recommend a pilot task-solving campaign be executed before the actual crowdsourcing project in order to determine how long the project will take to complete. Certain tasks are more difficult and complex than others, and

as such require more input and time from the crowdsourcers. The micropayments should thus be higher with such tasks, in order to achieve a comparable hourly wage. This was taken into account by e.g. Krek et al. (2013), proposing a micropayment of 0.02€ for simple crowdsourcing tasks (with approximately 350 decisions per hour, this amounts to 7€) and 0.05€ for more difficult tasks (the hourly wage remains the same, as the number of decisions per hour is somewhat smaller). When paying crowdsourcers the existing payment methods provisioned by local legislation and potential restrictions in tax legislation need to be taken into account.

Considering the ethical aspects of crowdsourcing is particularly important if the collected data will be used for commercial purposes and will be of direct financial benefit to the project initiator. In such cases it is controversial to offer crowdsourcers no or very low payment.

6.2 Recruitment restrictions

When selecting crowdsourcers for the project, potential legal restrictions need to be taken into account. This is especially important if minors are involved, and in most countries parental consent needs to be obtained before they can participate in a crowdsourcing project.

6.3 Authorship recognition

As crowdsourcers often do a significant amount of work on a project, it is necessary to determine how and where they will be credited. Although there are no clear guidelines when it comes to authorship recognition in crowdsourcing, several volunteer projects (e.g. *FoldIt*,¹² *Phylo*¹³) have listed crowdsourcers as project co-authors.

6.4 Intended use and distribution licence

Before starting work on a project, it is common for crowdsourcers to sign a consent form that informs them of the intended use of the collected data. The consent form needs to make clear whether their contribution will be used for research

¹² <https://fold.it/portal/>

¹³ <http://phylo.cs.mcgill.ca/>

purposes only or also for commercial ends, and whether it will be used in-house only or made available to third parties. If the crowdsourced data will be made publicly available, an adequate licence needs to be selected in accordance with local legislation on copyright and personal data protection.

7 CONCLUSION

Several projects related to the development of language resources and technologies have successfully implemented crowdsourcing, which indicates that this method could also prove useful in lexicography. However, all relevant aspects of this process need to be considered: from data preparation, microtask design and crowdsourcer recruitment, to ensuring crowdsourcer motivation and taking into account the legal, financial, and ethical restrictions of the project.

It is anticipated that crowdsourcing will soon be seen as a useful tool for lexicographers, one that will speed up the lexicographic process in a period of a growing demand for a the rapid processing of ever increasing amounts of linguistic data, as well as reduce the lexicographers' workload with regard to routine tasks, allowing them to focus on expert tasks. In addition, crowdsourced datasets will be useful for other purposes besides the construction of dictionaries, such as to help improve NLP tools through machine learning, with crowdsourced data as a high-quality training set. If adequately implemented, crowdsourcing could have a lasting impact on the workflow of future lexicographic projects, as well as on the use and life-cycle of lexicographic products.