Selected Crowdsourced Translation Practices

Authors: Aram Morera-Mesa, J.J. Collins, David Filip

Summary

This paper contains research related to workflow and design patterns. It briefly discusses the suitability of industry tools for crowdsourcing processes in terms of workflow pattern support. After listing a number of practices identified by analysing crowdsourced translation workflow models, the paper discusses four of the practices and presents two recommendations based on the scenarios of real crowdsourced translation projects.

Keywords: crowdsourcing, translation, localisation, crowdsourced translation, design patterns, workflow, workflow patterns

Background

This section presents the concepts of crowdsourcing, workflows, workflow patterns and design patterns and the relevant existing research.

Crowdsourcing

The introduction of crowdsourcing into localisation processes is a relatively new phenomenon. In this paper, crowdsourcing is understood as the practice of leveraging communities to carry out tasks that were traditionally carried out by professionals under contract (Howe 2006). Crowdsourcing translations as a way to deal with the increasing volume of content and locales has become a popular enough proposal within the industry that it has been argued that with a crowd of motivated, tech-savvy users, crowdsourcing is actually the best way of localising a product (Kelly 2009; Rickard 2009). Within academia crowdsourcing has been used in language related tasks, including translation, with an outcome of quality comparable to that of professionals (Callison-Burch 2009; Bloodgood and Callison-Burch 2010; Zaidan and Callison-Burch 2011).

The potential gain from adoption of this paradigm was illustrated by Facebook that in 2009 had sixty five languages available and an additional forty in production (DePalma and Kelly 2011). Although this success has lead industry experts to state that "*crowdsourcing will become integrated in the content supply chain*", most LSPs have not been able to find a way of integrating crowdsourcing in their processes and few other high profile cases have been reported (Désilets and van der Meer 2011).

Workflows

According to the Workflow Management Coalition a workflow is "*The automation of a business process, in whole or part, during which documents, information or tasks are passed from one participant to another for action, according to a set of procedural rules*" (WfMC 1999).

Business process models are similar, but simpler than what they represent (Maria 1997) since they do not represent information that is not relevant to the purpose of the model. This reduction of complexity makes the models useful for understanding and communicating the processes they represent, among other things (Giaglis 2001). There are a variety of business process modelling languages like Business Process Execution Language (BPEL), Yet Another Workflow Language (YAWL), XML Process Definition Language (XPDL) and many more (van der Aalst 2004) that can be used to define workflows and a number of systems able to enact them such as jBoss, Windows Workflow Foundation, WebSphere Process Server among others (Louridas 2008).

In 2011, Morera *et al* (2012) produced a series of workflow models for the processes of Asia Online's Wikipedia Translation Project as described by Vashee (2009), Facebook as described by Losse (2008), and Pootle and Crowdin. The models for Pootle and Crowdin were derived from use cases developed by the researcher. The use cases were created by carrying out multiple projects on both platforms simultaneously using several accounts in order to gain an understanding of the end-to-end process from the point of view of the different stakeholders.

Workflow Patterns

Workflow patterns are "series of constructs that are embodied in existing offers in response to actual modelling requirements" (Russell *et al* 2006). Control flow, data, resource, exception handling and presentation (van der Aalst *et al* 2011) are the currently recognized perspectives for workflow patterns, but in this paper the workflows are only analysed from a control flow perspective.

One of the applications of workflow patterns has been the evaluation of the expressive power of workflow modelling languages and execution platforms according to their support. In their research of 2011 Morera *et al*, concluded that out of 43 known control flow patterns, 13 were required to support crowdsourcing processes like the ones used for translation in the aforementioned projects and tools. After analysing the expressive possibilities of the modelling systems in two mainstream TMSs, the conclusion was that these only supported seven of the 13 required and were thus not suitable to automate that kind of crowdsourced translation process.

Design Patterns

Design patterns are solutions to re-occurring problems (Alexander *et al* 1977). Although Alexander *et al*'s work focused on architecture, the idea of design patterns extended to other fields like software architecture (Gamma *et al* 1995). There is an existing collection of design pattern for collaborative translation that was initially published by Désilets and van der Meer (2011) and has been further developed in a wiki at collaborative-translation-patterns.com. The initial patterns emerged from a one day workshop with expert and covered many aspects of crowdsourced translation projects.

New and updated design patterns

The patterns listed in this paper emerged from the analysis of the workflow models of the processes used or enabled by Crowdin, Facebook, Asia Online Wikipedia translation project, DotSub, Amara, Kiva, Launchpad's Rosetta and Pootle and processes described in the literature, specifically the work of Exton *et al* (2009) and Zaidan and Callison-Burch (2011) for hidden redundant assessment and Désilets and van der Meer (2011) for content selection and unit granularity selection. In all cases the criterion to add a pattern to the list was its appearing in more than one process.

In order to refine the patterns eight interviews with experts were carried out. As visible on table one, the interviews were carried out in the months of April and May of 2013 and lasted between 47 and 81 minutes.

Interviewee	Date	Duration
Subject P	05/04/2013	49:20
Subject C	11/04/2013	1:11:02
Subject R	17/04/2013	41:41
Subject M	18/04/2013	57:15
Subject H	02/05/2013	1:11:47
Subject S	17/05/2013	47:31
Subject A	20/05/2013	1:20:59
Subject V	21/05/2013	57:31

The interviews yielded 64.456 words of transcripts that were analysed using thematic analysis which was chosen because of the approach allowing the use of top-down, deductive themes that emerged from prior research and themes generated inductively from the data (Boyatzis 1998, p.4) and its combination of views from both the interpretivist and positivist paradigms (Guest *et al* 2011, p.15) that matched the pragmatic approach that the researcher wanted to follow.

Chart one shows the 16 patterns and the number utterances related to each of them in the interviews.
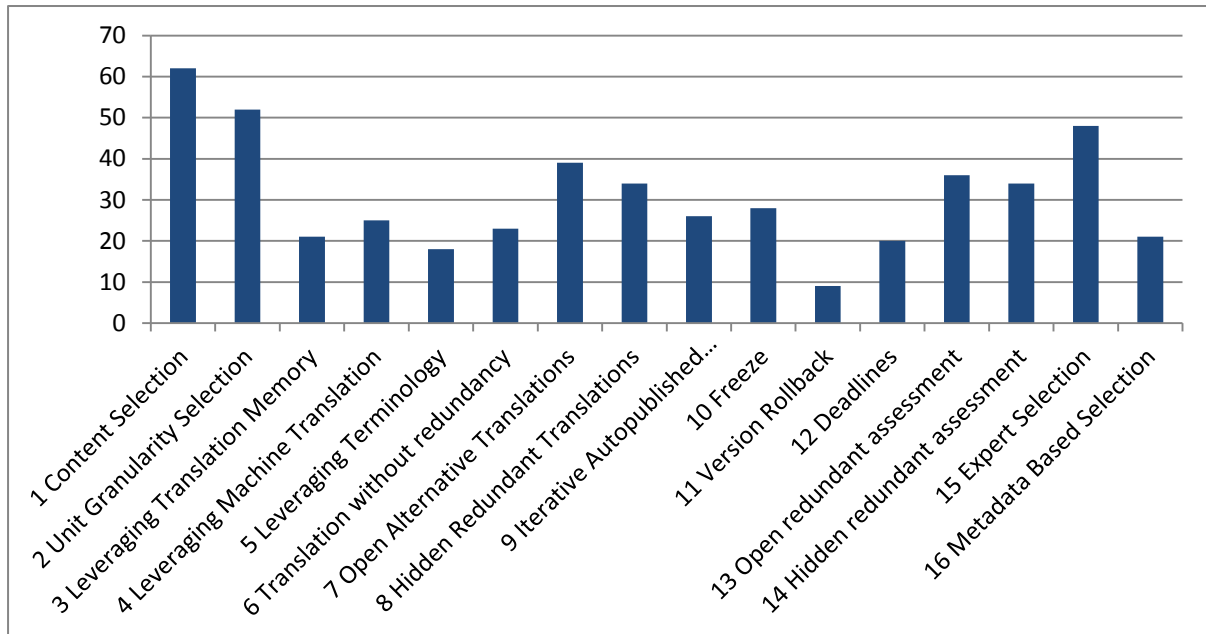

Chart 1

Selected Patterns

This section presents four of the patterns together with the conclusions reached from the statements from the interviewees.

Pattern 7: Open Alternative Translations

This is the practice of collecting multiple translations for a single source Translations Unit (TU), i.e. the smallest amount of text that is processed as a single unit. Translation units are often sentences, but can also be strings, paragraphs or even whole texts, depending on the project. When this pattern is implemented collaborators are able to see other existing translations. Facebook, Crowdin and Launchpad's Rosetta use this pattern.

Key findings

According to the experts interviewed, using open alternative translations:

- Enables the open and hidden crowd assessment patterns.
- Fosters crowd engagement.
- Wastes effort in the creation of translations that will not be used.
- Generates noise that affects the selection process.

- Generates redundant data that weights the databases supporting them.

- Can result in a challenging process to select the best translation.

- Enables Expert Selection and Review pattern, where the crowd suggestions can be used as a starting point for the work of a professional or trusted member of the community.

- Produces variations of translations that can be used to train statistical MT systems.

- Is better suited for working with smaller TUs.

Open Alternative Translations discussion

There were two features of this pattern that the researcher expected the experts to bring up but were not discussed by any of the interviewees. First is the effect on contributors of seeing existing translations.

Existing translations can work to some degree as TM and MT for contributors in that the contributors can copy and paste an existing translation and modify those parts upon which they feel they can improve. This would accelerate the process for the creation of individual new suggestions.

Furthermore, in the same way that Subject P worries about novice translators imitating the style of an MT system, the researcher thinks that existing translations will impact the following translations. For example, if all existing translations use a given term even without an authoritative translation for it, it is unlikely that a new alternative will decide to use a different term.

This would be an example of the convergence processes discussed by Lorenz *et al* (2011). The convergence processes are:

The 'social influence effect', which says that "social influence diminishes diversity in groups without improving its accuracy" (*ibidem*).

The 'range reduction effect' which says that the position of the truth, in this case that would be the optimal translation, over time moves to peripheral regions (*ibidem*), that is, becomes less similar to the cluster of translations that are the most popular.

The 'confidence effect' through which "opinion convergence boosts individuals' confidence in their estimates despite a lack of collective improvements in accuracy" (*ibidem*).

The other feature of this approach that was not discussed is that material rewards for contributors are less helpful if applying this practice. They either would have to be too small to actually help motivation in order for the cost not to become prohibitive by accumulation of numerous alternatives for each source TU or would have to be limited to the selected translation. Rewarding only the

translations that are selected could lead to some contributors focusing on being selected and harm the community dynamics.

Pattern 8: Hidden redundant translations

This is the practice of collecting multiple, limited in number, alternative translations for a single TU without letting the contributors see the other translations.

On the industry side, txtEagle and CrowdFlower collect multiple redundant translations, but do not display them to the users (Eagle 2009; Bentivogli *et al* 2011), in academia this approach has also been successfully used through AMT (Zaidan and Callison-Burch 2011). In all three cases, translators obtain material rewards for their contributions. This pattern was also enacted in Subject A's project where new, untrusted volunteers did not see the suggestions of other volunteers.

Many of the features of unlimited redundant alternatives apply, and the differences will be covered in the discussion part for this pattern.

Key findings

According to the experts interviewed, using open alternative translations:

- Enables frequency based automated selection of translation.
- When combined with metadata based selection, increases the throughput by directing the effort towards those TU that are yet to be translated.
- When compared to using Open Alternative Translations, improves the spread of effort among volunteers by directing them to TUs that do not sufficient translations yet.
- If an expert selects the final translation, saves time by simplifying the decision process.
- If a paid expert selects the final translation, saves cost by limiting the number of words that the expert has to consider.
- Can potentially produce more variants that are useful for MT training.
- Prevents you from using community activities such as voting and commenting on the translations. This will damage the community related motivations of your contributors.
- May result in your needing material incentives in order to replace the community incentives.
- Run the risks that translators feel like they are not being appreciated and forces you to mitigate this risk in your communications strategy.
- You may find out that the optimal translation has not been suggested and the quality of the published translation ends up being lower than when using other approaches
- May create redundant translations that do not add value because a better translation already exists but your contributors are not aware of it because it is hidden.

- Enables you to combine redundancy with longer TUs, but this would enter in conflict with having a communication strategy that addresses what happens with the translations.

Hidden Alternative Translations discussion

Most points considered by the researcher were discussed by the interviewees. However, there were some features that none of the interviewees brought up.

First, by collecting a limited number of translations the generation of translations itself would take less time since, for example, it takes less time to collect five translations than to collect thirty.

Second, this practice by setting limits to the cost enables you to use redundancy even if you are offering material rewards for the translations. This is illustrated by the processes used by TxtEagle (Eagle 2009), Crowdflower (Bentivogli *et al* 2011) and Zaidan and Callison-Burch (2011).

Third, by hiding existing translations the contributors are freed from their influence hence avoiding the 'social influence effect' (Lorenz *et al* 2011). Subject A suggested that hidden alternatives produces more variants. The increased number of variants could be consequence of having avoided the social influence effect.

Also, as discussed in the previous pattern, the researcher thinks that novice translators may imitate existing translations in a project using the open alternative translation, but this effect would be avoided using hidden alternatives.

Lastly, although the researcher agrees with one of the experts that using hidden alternative translations with longer TUs is feasible, the researcher thinks that using redundancy with longer work units, unless the contributors are being motivated by material rewards, will generally produce bad results because of the issues related to the Unit granularity selection pattern, which is not discussed in this paper but is part of the researcher's thesis.

Pattern 13: Open redundant assessment

This pattern is enabled for example by Facebook. This is usually enacted by letting contributors up or down-vote translations. Another type of assessment seeing in the literature is ranking, but that has only been used in hidden assessment (Zaidan and Callison-Burch 2011).

Key findings

According to the experts interviewed, using open redundant assessment:

- Enables vote based selection, which results in a more crowd-managed process.
- Crowd assessment is more trustworthy for larger crowds.
- Crowd assessment adds a level of transparency that can accelerate the process by allowing translation efforts to be directed to untranslated text.
- By making the process more self-managed, working with unlimited redundant alternatives becomes less challenging.
- Can work well if the criterion for quality is the acceptability of a translation; however:
  - If no measures are taken to prevent it, voting can cause a Yule process that hides better translations than the ones currently used, but that does not mean the ones selected are not good enough.
- Can be combined with expert selection in two ways:
  - Crowd assessment can be used to pre-select translations before using the expert selection pattern to select the published one
  - Expert selection could be used to preselect translations before letting the crowd choose the published one.
- Crowd assessment opens another venue for crowd engagement, which is valuable for marketing purposes.
- Crowd assessment is open to vandalism.
- To prevent vandalism, votes should have provenance metadata and be weighted according to the voter.
- There must be a plan in place to handle the controversies that may arise with crowd assessment.
- Having a specific UI for assessment is desirable.
- Crowd assessment for specialized documents is only suitable if the crowd is made of specialists too.
- Crowd assessment is not suitable if the intention is to have translators develop a personal style.
- Crowd assessment is not suitable for languages that do not have a translation tradition, because the crowd will probably agree on a less than optimal translation.
- Crowd assessment is not suitable for longer TUs.

Open redundant assessment discussion

The coverage for this pattern went beyond the aspects foreseen by the researcher; this may be due to the pattern being implemented by Facebook and several of the interviewees. However, none of the interviewees considered forms of assessment other than voting. As discussed before, Zaidan and Callison-Burch (2011) successfully used ranking of translations in their experiment. Ranking requires a higher cognitive effort than voting, so it may not be suitable for many crowdsourcing contexts, however, if the crowd is being paid, as it was the case for Zaidan and Callison-Burch, ranking is another type of open or hidden assessment that could be implemented.

Regarding the Yule process risk caused by visibility, the experiment performed by Muchnik *et al* (2013) showed that positive votes feed positive votes creating bubbles while negative votes tend to be neutralized over time. This effect would further accentuate the Yule process, so it would be recommendable to take some measures to reduce it if the criteria for quality were not solely the popularity of a translation.

Pattern 16: Metadata based selection

Metadata based selection is the practice of automatically choosing a translation from among multiple alternatives by using the metadata attached to it. There are four approaches to it in use: Trusted translator, time stamp, frequency and crowd assessment. In this paper only frequency and crowd assessment are addressed.

Key findings

- Automated metadata based selection can be used with:
    - Frequency of translations if you have used hidden redundant alternatives.
    - Votes both in the form of open or hidden redundant assessment.
- Automated selection is the fastest selection method.
- Automated selection can be the only viable solution in extremely big projects.
- Automated selection has consistent criteria.
- Automated selection can be tricked, but some mechanisms to minimize this risk can be put in place.
- Automated selection should only be used where there is trust in its fitness for the purpose.
- Automated selection should not be used for controversial or sensitive content.

Metadata based selection discussion

In the mind of the researchers, using votes in the form of open assessment is by far the most transparent approach. This transparency can make the crowd feel empowered and be beneficial for their motivation. However, as discussed before, the social influence effect may prevent the best translations from being chosen when using this pattern.

Using hidden votes should free the voters from the social influence effect, but the loss in transparency may hurt the motivation of the crowd.

Using frequency as a criterion is also lacking in transparency, but by skipping the need to have a voting stage, it is potentially the most time efficient approach of the three approaches considered in this paper.

Patterns in context

Facebook case

Model one represents the process described by Losse (2008) for community supported languages in Facebook. The purpose of enabling the users to translate Facebook was having the platform available in as many languages as possible (*ibidem*). Users can translate the UI and documentation, but not legal documents, as the result of Facebook **selecting suitable content**. In this process a user can see the available TUs, which are strings, one of the possible **TU granularity selections**, and decide to vote for one translation, update their vote to give it to a different translation or make a new suggestion. If they decide to add a new suggestion, Facebook displays relevant **terminology**. Unless the **freeze** task that appears towards the end of the process has been executed, the possibility of voting, updating the vote or suggesting new translations is always open to other contributors. Since all contributors can see existing votes and translations, it can be said that this stage implements both **open alternative translations** and **open crowd assessment** patterns. In this process the translation is selected automatically according to the number of votes, which is an instantiation of the **metadata based selection** (open crowd assessment) pattern.
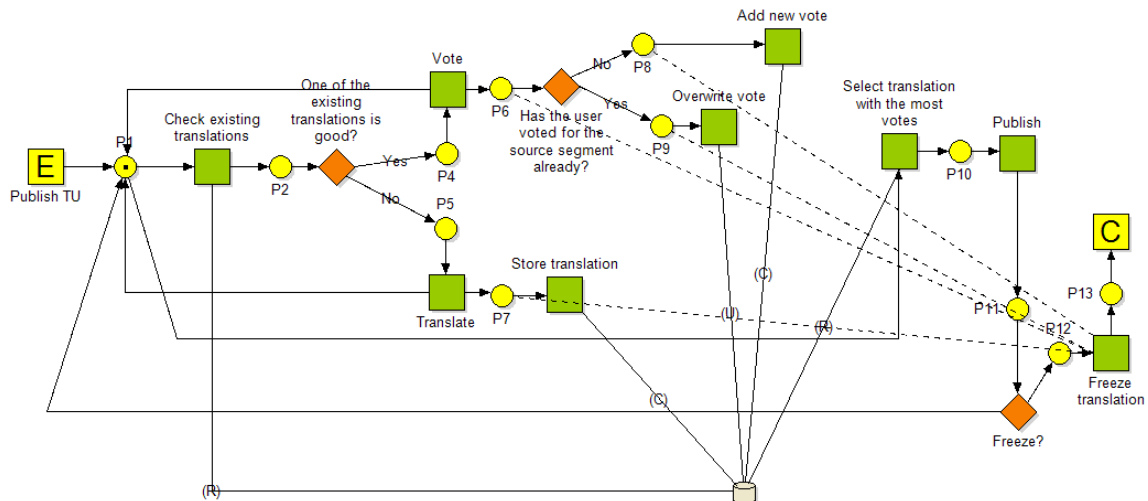
Model 1

Asia Online Wikipedia translation case

Model two represents the process described by Vashee (2009) for the Asia Online Wikipedia translation project. The purpose of this project was to translate billions of words to increase Thai internet content (DePalma 2011). The **selected suitable content** was the English Wikipedia. In this project, contributors were presented with a collection of sentences resulting from having segmented Wikipedia entries, this is an implementation of **TU granularity selection** (segmented longer unit). In this process three contributors post edit the output of an **MT** system. The contributors cannot see the translations of the others, making this an implementation of the **hidden alternative translations** pattern. The three postedited translations are then automatically compared. If two of them match, the third translation is discarded and the matching translation is stored for MT training and published. This is an example of **metadata based selection** (frequency). If none of them match a professional selects one of them and decides if the other two are good translations that will be kept in a TM for MT training or are discarded. This is an example of **expert selection.**
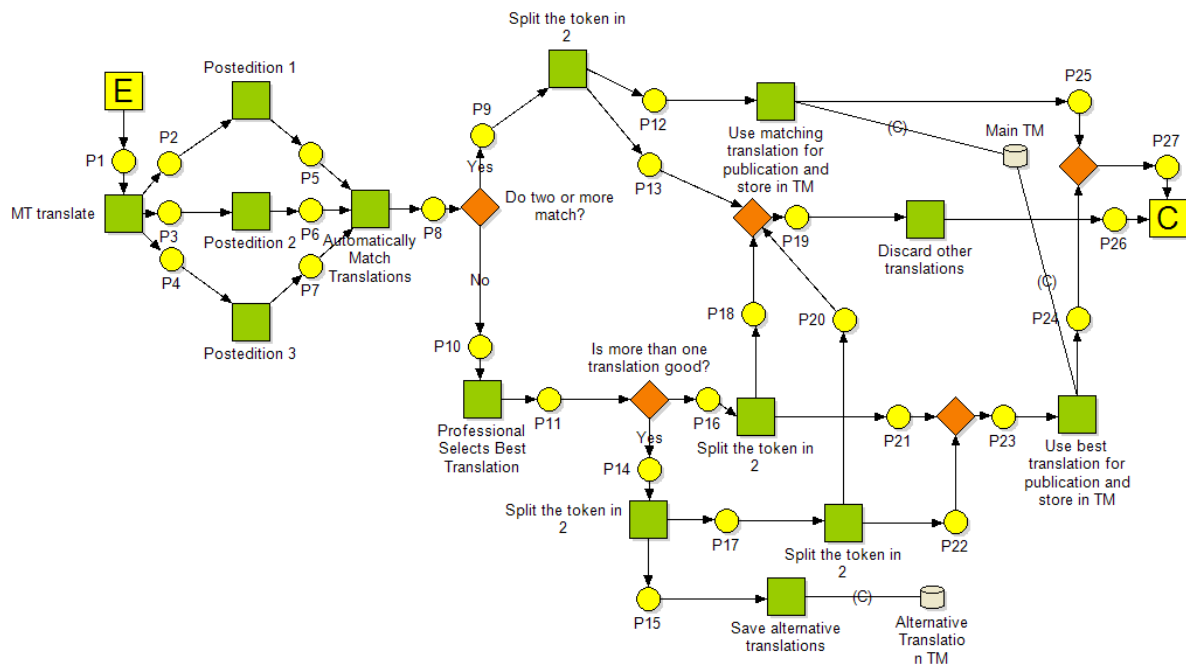
Model 2

Conclusion

This paper has covered the concepts of crowdsourcing, workflows, workflow patterns and design patterns. It has introduced a collection of new or updated design patterns for crowdsourced translation and presented four of them in more depth. Finally the paper has presented two workflow models and pointed the design patterns that they implement.

The two examples illustrate two paradigms of crowdsourced translations. In the case of Asia Online's Wikipedia translation project, crowdsourcing was used to reach the highest possible throughput while avoiding translations with typical MT output issues. In the views of the researcher, projects like this would do well following their example and implement the following patterns:

- TU granularity selection (segmented longer unit) pattern, that allows contributors to translate small units while still having contextual information necessary to make sense out of them.

- Leverage MT to accelerate the work of the contributors and potentially include the right terminology in the output if the right training material was used.

- Hidden alternative translations, to enable automatic metadata based selection (frequency) and avoid excessive redundancy.

- Metadata based selection (frequency) for the fastest possible selection that is not submitted to the social influence effect.

Although in her talk Losse said that the objective of having Facebook translated by the users was to reach as many locales as possible, their approach is also very suitable for increasing customer engagement and brand loyal, which is one of the main reasons for organizations to use crowdsourcing (Désilets and van der Meer 2011). In that context the organization will want the most transparent process possible and the most power possible in the hands of the contributors. Bearing this in mind the organization should use:

- TU granularity selection (segments or strings) pattern, since the more units you have the more room there is for people to participate.
- Open alternative translations, to allow for a virtually unlimited number of translation being suggested, providing hence more room for contributors to participate.
- Open crowd assessment, which again allow for a virtually unlimited number of contributors to get involved, plus allows them to see the assessments of other contributors resulting in more transparency.
- Metadata based selection (open crowd assessment), which conveys to the crowd the idea that the finally selected translation is in their hands.

There are other types of crowdsourced translation that do not use any kind of redundancy or work in a more iterative way, use rollbacks, different TU sizes, etc. and potentially more patterns for the generation of translation payload and metadata. There are also other aspects of crowdsourced translations that deserve further research in the collection of best practices. This paper is only a second attempt and the authors are excited to see the development of more complete collections of patterns for crowdsourced translation.

Bibliography

Van der Aalst, W.M.P. (2004) 'Business process management demystified: A tutorial on models, systems and standards for workflow management', *Lectures on Concurrency and Petri Nets*, 21–58.

Van der Aalst, W.M.P., Ter Hofstede, A.H.M., Russell, N. (2011) Workflow Patterns [online], available: http://www.workflowpatterns.com/.

Alexander, C., Ishikawa, S., Silverstein, M. (1977) *A Pattern Language: Towns, Buildings, Construction*, Oxford University Press, USA.

Bentivogli, L., Federico, M., Moretti, G., Paul, M. (2011) 'Getting expert quality from the crowd for machine translation evaluation', *Proceedings of the MT Summmit*, 13, 521–528.

Bloodgood, M., Callison-Burch, C. (2010) 'Using Mechanical Turk to build machine translation evaluation sets', Association for Computational Linguistics, 208–211.

Boyatzis, R.E. (1998) *Transforming Qualitative Information: Thematic Analysis and Code Development*, Sage.

Callison-Burch, C. (2009) 'Fast, cheap, and creative: evaluating translation quality using Amazon's Mechanical Turk', Presented at the Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1-Volume 1, Association for Computational Linguistics, 286–295.

DePalma, D. a. (2011) 'The Largest Translation Project…So Far', available: http://www.commonsenseadvisory.com/Default.aspx?Contenttype=ArticleDetAD&tabID=63&Aid=1180&moduleId=390.

DePalma, D.A., Kelly, N. (2011) 'Project management for crowdsourced translation', *Translation and Localization Project Management: The Art of the Possible*, 379.

Désilets, A., van der Meer, R. (2011) 'Co-creating a repository of best-practices for collaborative translators', *Linguistica Antverpiensia*, 10.

Eagle, N. (2009) 'txteagle: Mobile crowdsourcing', *Internationalization, Design and Global Development*, 447–456.

Exton, C., Wasala, A., Buckley, J., Schäler, R. (2009) 'Micro Crowdsourcing: A new Model for Software Localisation', *Localisation Focus*, 81.

Gamma, E., Helm, R., Johnson, R., Vlissides, J. (1995) 'Design patterns. Elements of reusable object-oriented software', *Addison-Wesley Professional Computing Series, Reading, Mass.: Addison-Wesley,| c1995*, 1.

Giaglis, G.M. (2001) 'A taxonomy of business process modeling and information systems modeling techniques', *International Journal of Flexible Manufacturing Systems*, 13(2), 209–228.

Guest, G., MacQueen, K.M., Namey, E.E. (2011) *Applied Thematic Analysis*, Sage.

Howe, J. (2006) 'Crowdsourcing: A definition', *Crowdsourcing: Tracking the rise of the amateur*.

Kelly, N. (2009) Freelance Translators Clash with LinkedIn over Crowdsourced Translation [online], available: http://www.commonsenseadvisory.com/Default.aspx?Contenttype=ArticleDetAD&tabID=63&Aid=591&moduleId=391.

Lorenz, J., Rauhut, H., Schweitzer, F., Helbing, D. (2011) 'How social influence can undermine the wisdom of crowd effect', *Proceedings of the National Academy of Sciences*, 108(22), 9020–9025.

Losse, K. (2008) *Facebook - Achieving Quality in a Crowd-sourced Translation Environment* [online], available: http://www.localisation.ie/resources/presentations/videos/video2.htm.

Louridas, P. (2008) 'Orchestrating web services with bpel', *IEEE Software*, 25(2), 85–87.

Maria, A. (1997) 'Introduction to modeling and simulation', Presented at the Proceedings of the 29th conference on Winter simulation, IEEE Computer Society, 7–13.

Morera, A., Aouad, L., Collins, J. (2012) 'Assessing Support for Community Workflows in Localisation', Presented at the Business Process Management Workshops, Springer, 195–206.

Muchnik, L., Aral, S., Taylor, S. (2013) 'Social Influence Bias: A Randomized Experiment', *Science*, 341(6146), 647–651.

Rickard, J. (2009) 'Translation in the Community', in *LRC XIV Localisation in The Cloud*.

Russell, N., Hofstede, A.H.M., Aalst, W.M.P., Mulyar, N. (2006) *Workflow Control-Flow Patterns: A Revised View, BPM Center Report*, BPM-06-22, BPMcenter. org.

Vashee, K. (2009) 'MT Technology in the Cloud – An evolving model', in *LRC XIV, Localisation in The Cloud*, available: http://www.localisation.ie/resources/conferences/2009/presentations/LRC09-KV.pdf.

WfMC (1999) *Workflow Management Coalition: Terminology & Glossary* [online], available: http://www.wfmc.org/standards/docs/TC-1011_term_glossary_v3.pdf.

Zaidan, O.F., Callison-Burch, C. (2011) 'Crowdsourcing translation: professional quality from non-professionals', *Proceedings of ACL 2011*.

Biographical information

Aram Morera Mesa is a PhD researcher at University of Limerick. He obtained his BA in Translation and Interpreting by Universidad Complutense de Madrid in 2005. Between 2005 and 2008 he worked as a freelance translator and teaching assistant. In 2009 he obtained a Graduate Dimploma in Localisation by the University of Limerick, where he has been carrying out research in localisation up to the present day. He can be contacted at: aram.morera-mesa@ul.ie

J.J. Collins is a Lecturer in the Dept. of Computer Science of the University of Limerick. He holds a MSc in Computer Sciences, Artificial Intelligence stream by the University of London. His research interests include, among others, artificial intelligence, software architecture and design patterns. He can be contacted at: j.j.collins@ul.ie

David Filip is the Liaison Officer, Secretary and Editor of OASIS XLIFF TC, Co-Chair of W3C MultilingualWeb-LT (ITS 2.0), XLIFF TC Liaison at Unicode Localization Interoperability (ULI) TC, TAUS Standards Advisory Board Member etc. His specialties include open standards and process metadata, workflow and meta-workflow automation. David works as a Research Fellow at the Localisation Research Centre (LRC) at the University of Limerick, Ireland in the capacity of a CNGL Challenge Leader in the area of Open Standards for Global Intelligent Content (GLOBIC) Interoperability. He has been engaged in a number of publicly funded research and development projects such as the Centre for Next Generation Localisation (CNGL both Mark I Mark II), LT-Web, and LetsMT!. Before 2011, David oversaw key research and change projects for Moravia's worldwide operations. He has held research scholarships at universities in Vienna, Hamburg and Geneva, and he graduated in 2004 from Brno University with a PhD in analytic philosophy. David also holds master's degrees in philosophy, art history, theory of art and German philology. His academic theses dealt with the practical application of analytic methodologies, formal semantics and translatability. He can be contacted at: david.filip@ul.ie