

LingoTowns: A Virtual World For Natural Language Annotation and Language Learning

CHRIS MADGE, Queen Mary University Of London, UK
JUSSI BRIGHTMORE, Queen Mary University Of London, UK
DORUK KICIKOGLU, Queen Mary University Of London, UK
FATIMA ALTHANI, Queen Mary University Of London, UK
RICHARD BARTLE, University Of Essex, UK
JON CHAMBERLAIN, University Of Essex, UK
UDO KRUSCHWITZ, University of Regensburg, DE
MASSIMO POESIO, Queen Mary University Of London, UK

In this paper we introduce *LingoTowns*, a new GWAP platform targeting language learners. *LingoTowns* provides a unified experience integrating games for multiple aspects of lexical and grammatical experience in a single virtual world, whilst simultaneously collecting judgements. Both *LingoTowns* and its constituent games are designed to provide more engagement to the players/ learners than normal GWAPs. The platform also incorporates knowledge tracing methods ensuring that the players' progress in terms of understanding of grammatical concepts is tracked both at the individual game level and overall.

CCS Concepts: • **Applied computing** → *Interactive learning environments; Computer games*; • **Human-centered computing** → *Natural language interfaces; Web-based interaction*; • **Software and its engineering** → **Interactive games**.

Additional Key Words and Phrases: Games With A Purpose, Educational Games, Virtual Worlds

ACM Reference Format:

Chris Madge, Jussi Brightmore, Doruk Kicikoglu, Fatima Althani, Richard Bartle, Jon Chamberlain, Udo Kruschwitz, and Massimo Poesio. 2022. LingoTowns: A Virtual World For Natural Language Annotation and Language Learning. In *Extended Abstracts of the Annual Symposium on Computer-Human Interaction in Play (CHI PLAY '22 EA)*, November 2–5, 2022, Bremen, Germany. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3505270.3558323>

1 INTRODUCTION

Games with a Purpose (GWAPs) [16, 27] for different types of data labelling have been under development for over twenty years, but while GWAPs in some domains have fulfilled their promise of leveraging the web population to label vast amounts of data at low cost (one obvious example of successful GWAP being *FoldIt* [9]), others have proved more challenging. NLP, in particular, is an area in which GWAPs have exceeded traditional crowdsourcing methods and demonstrated great potential [15, 22], but are yet to rival the success seen in other domains.

One key prerequisite for the success of an online annotation platform or GWAP is participant's **motivation**. Surveys of contributors to *GalaxyZoo* [23] or *FoldIt* [9] reveal the desire to make a scientific contribution as the primary motivation,

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2022 Copyright held by the owner/author(s).

Manuscript submitted to ACM

but not many GWAPs have been successful at motivating contributors this way. Engaging players in entertaining activities has been the key motivation that GWAPs attempt to achieve [27], but not many actually deployed GWAPs for NLP are truly entertaining. For example, our own *Phrase Detectives* [22] has been moderately successful, attracting tens of thousands of players and managing to collect over five million annotations at half the estimated cost in comparison with micro-task crowdsourcing approaches [22]. However, despite all the gamification components that do motivate play, ultimately one of the core motivators was still financial incentives, in the form of small prizes for the players [5].

Thus, a number of more recent games, most notably *DuoLingo* [28], have been using **language learning** as motivation for their players.¹ *Phrase Detectives* itself has been used in university-level language courses since 2011. There is to be a natural synergy between the objectives of the developers of GWAPs for NLP and the interests of language teachers and language learners using language learning apps. Learners wanting to improve their knowledge of a language would appear to be excellent targets for GWAPs; and conversely, GWAPs teaching grammar concepts to their players could be good ways to practice grammar, provided they are sufficiently entertaining, teach the necessary concepts clearly, and provide useful feedback. Learning has been said to be part of what makes a game fun [14]. Conversely, teaching the players about grammatical concepts also promises to improve the quality of their future contributions. Leveraging games' natural ability (through tutorials etc.) to teach players in this way has been said to be one of the key benefits provided by GWAPs over other crowdsourcing methods [24].

In this demo paper we present *LingoTowns* (demonstration video available ²), a new GWAP platform for smartphones that targets language learners but also addresses other limitations we have come to recognize in fifteen years of work with *Phrase Detectives*.

First of all, *LingoTowns* integrates games for many aspects of language interpretation, unlike most NLP GWAPs. From an NLP perspective, **gamifying the entire pipeline** could improve the quality of the collected data. NLP annotations typically require annotating multiple interdependent levels. For example, annotating coreference requires identification of the mentions in a text, which in turn requires syntactic annotation or at least part-of-speech annotation, which in turn requires identifying token boundaries. By including in the platform games for all these levels we can ensure that the input to each game had been verified by the players of the previous game in the pipeline. From a language learning perspective, such a platform would make it possible for grammar learners to learn and be evaluated about different aspects of grammar. In *LingoTowns*, multiple GWAPs, each with their own features and interface, designed specifically to suit their respective annotation tasks, are all grouped under one platform. To our knowledge, only one platform of this type has been attempted in NLP, *WordRobe* [26], whose design however substantially differs from that of *LingoTowns*, as discussed below.

Our second key objective was to make the platform as **entertaining and game-like** as possible. We developed distinct games for each NLP task, each meant to be a proper game as opposed to a gamified annotation activity. For *LingoTowns* as a whole, we looked for an overarching game genre that would not break player immersion; this provides a lot of freedom and flexibility for the designer of the mini-game. We settled on the **virtual word** [3] paradigm, which provides the opportunity to inject a sense of exploration, fantasy and agency into the design of GWAPs. We also used **procedural content generation** to ensure the players would have access to infinite content.

¹Although *DuoLingo* was originally conceived as a GWAP to collect data for machine translation, that objective has now been abandoned and the focus is now solely on developing a language learning app.

²<https://www.youtube.com/watch?v=d3a0i1RkYv0>

One of the key reasons for the success of *FoldIt*, but also of *Phrase Detectives*, is its **usability**, which is crucial to engagement as poor usability leads to poor flow. *LingoTowns* was the result of a thorough re-analysis and re-design of our existing GWAPS based on user input.

The third key objective of *LingoTowns* is to incorporate both **progression** and **task assignment** strategies. In *Phrase Detectives*, all players that pass an initial test are then assumed to be able to carry out the same labelling tasks, which preliminary tests suggested could be carried out with limited annotator disagreement. Our experience however showed that the result was that some players were asked to express judgments that were too complex for them (e.g., marking plurals) whereas other, linguistically trained players, were not allowed to provide more complex judgments. In *LingoTowns* **knowledge tracing** methods [8] are used to ensure that players are continuously assessed on their understanding of the key linguistic concepts, and may eventually progress to more complex tasks.

Effective task assignment strategies are a very important element of human computation systems. However, these strategies do not necessarily apply to GWAPS. Besides accuracy, GWAPS also target entertainment and the same progressive experience used by games [19]. Previous work has explored using game based methods of assigning tasks to players in joint pursuit of entertainment and accuracy through adaptation of Elo and Glicko2 to player-vs-task rather than player-vs-player [7]. In this work, we centralise the progressive task-assignment between multiple tasks. This approach provides the opportunity to explore methods of progression between games and annotation tasks that allows us to strike a balance between delivering: an entertaining experience; an educational experience and an optimised strategy for the final accuracy of the annotations.

2 BACKGROUND

Games with a Purpose (GWAPS) seek to gather annotations as a byproduct of play [27], thereby motivating the player with entertainment rather than a financial incentive. Originally, these games targeted image labelling tasks [29], but this approach was soon applied to gathering NLP labels. One example is *Phrase Detectives* [22], that gathers labels for coreference. Other examples include games targeting mention detection [18], semantic relations [15], part-of-speech tagging [17] and many others.

To our knowledge, there have been three previous examples of **virtual world-based GWAP** role-playing games. *High School Superhero*, A 3D abusive language detection game [4]; *Katana and Grand Guru: a Game of the Lost Words* [21]; and *The Knowledge Towers* [25]. However, these games are used directly to gather a single type of annotation, rather than serve as a meta-game for multiple optionally independent annotation GWAPS. In traditional game design, the concept of wrapping a series of mini-games in a master game that supports progression between tasks is commonplace. For example, the virtual world games *Pokémon*; *Minecraft*; *Animal Crossing*; *Final Fantasy XV*; *The Legend Of Zelda* and *Red Dead Redemption 2* amongst others, all incorporate a fishing mini-game, whilst not being fishing orientated games.

There are, to our knowledge, only one example of a **multi-task GWAP**. The *LingoBoingo* portal³ provides access to a collection of NLP games in a single portal. However, the portal itself is not gamified. To our knowledge, *Wordrobe* is the only example of a game addressing multiple truly distinct annotation tasks, cleverly wrapped with one game mechanic [26]. However, this clever design is not without limitation. *Wordrobe* gives a continuous experience throughout the game by presenting all annotation tasks as multiple choice questions. This interface lacks the interactivity required for users to completely express their annotation choices. Instead, they must select one of several possibly correct options.

³<https://lingoboingo.org/>

The application **gamification in language learning** is now a very popular approach. Examples include Memrise ⁴, with a focus on real-life language learning using video clips from native speakers, and Babbel ⁵, that features live lessons and podcasts. Of these, Duolingo ⁶ is one of the inspirations for our approach. Originally sporting the tag line, “Learn a language, translate the web”, Duolingo originally offered learning while gathering annotations.

This **trade-off between reading and labelling** is a focus in crowdsourcing, particularly in active learning settings where the goal is to minimise the number of annotations required [20, 30]. The idea being, that reading and understanding, like labelling, is time consuming and therefore costly, so it makes sense, in the interest of resource optimisation, having familiarised an annotator with a document to have them continue to annotate that document or provide multiple types of labels on the same document. In this work, we gamify this approach, by having a player provide multiple types of annotation over the same document.

3 DESIGN

In this section we will describe the design of the *LingoTowns* game. Broadly speaking, *LingoTowns* is a web based, infinite, procedurally generated isometric world. Each town is associated with a single document, and each building in that town represents a single annotation activity/game.

3.1 Story

In *LingoTowns*, players find themselves in a future Earth in which humans have forgotten their rich history. Surviving documents have been unearthed in towns across the world, but the evolution of human language means they are no longer understood. The future civilisation transports the player forward in time, to help them decode the documents and restart the towns.

3.2 Aesthetics

The experience is presented as an isometric view of a world, as in Figure 1. The world layout is coupled with the task assignment. The player is motivated to complete tasks partly through the freedom it gives them to continue to explore the world. As a consequence, the world is “infinite”, locations can be dynamically added in-line with the players progression and the requirement for annotations.

The available space is split into levels. At any one time, a player may choose from one or more document/task sets, selected for them by the task recommendation algorithm that are suitable for their assessed ability level. This provides the player with an important sense of agency.

The levels are laid out in concentric circles on the map, from the “home position” at the very centre of the map, moving outwards with increasing difficulty (Figure 3). The player’s view of future tasks is obscured by a cloud covering (shown in Figure 1), that provides a horizon beyond which the player cannot see, until the system has selected a new set of tasks for them. The system selects future tasks based on their previous behaviour. Having selected the next set of available tasks, new towns (one for each document) are created in the next circle with a random radial placement, and the cloud cover recedes to allow the player to visit them. Each town has buildings that represent the games available to play in them. To give some additional distinction over the separation of towns, each town has a theme/biome (e.g. Farmland, Woodland, Lakes, Desert).

⁴<https://memrise.com>

⁵<https://babbel.com>

⁶<https://www.duolingo.com>

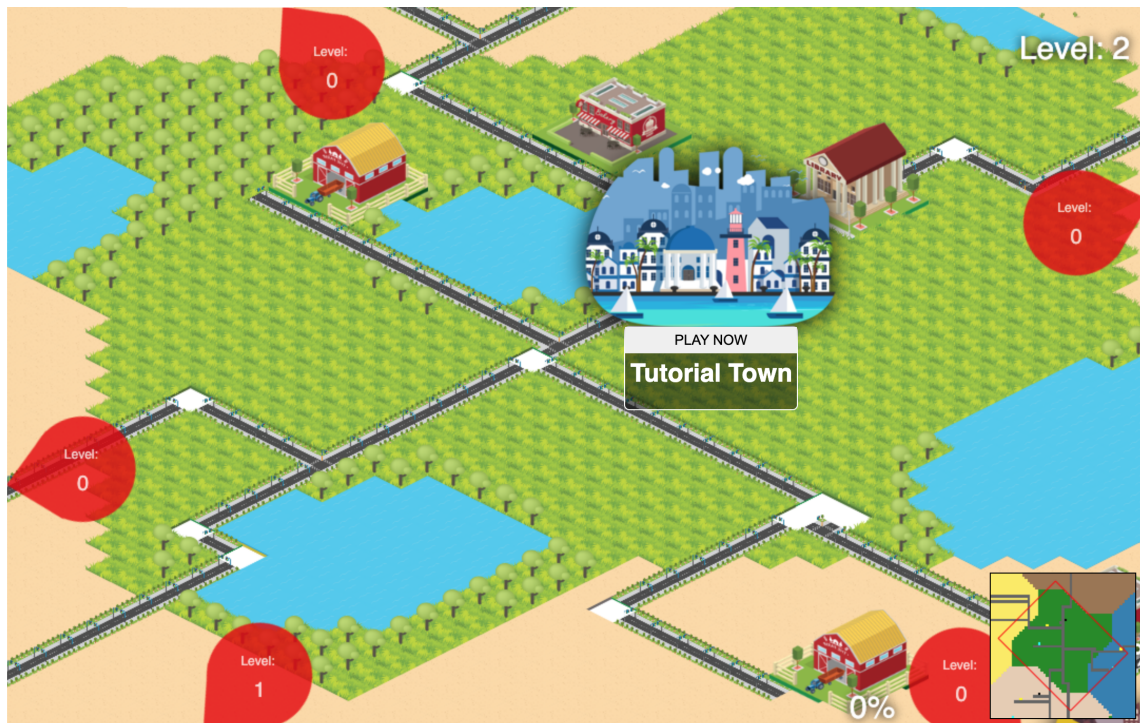


Fig. 1. The *LingoTowns* World Map

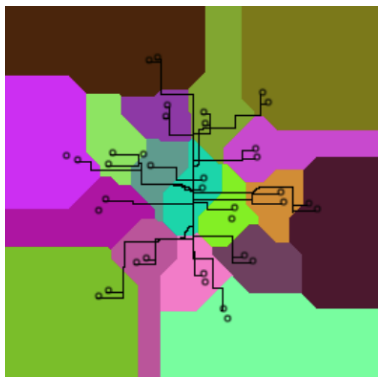


Fig. 2. Layout Generation (biomes: coloured regions; roads: lines; buildings: circles)




Fig. 3. Game World Layout

3.3 Generation

Each player sees a unique world tailored to their learning experience and progression that continues to evolve with their play. To achieve this whilst meeting the aforementioned design, procedural generation is used.

Each time a new level of towns is unlocked, the towns are placed equidistant in a radial configuration. Each towns buildings are then placed using Poisson Disc sampling [6]. This produces a random layout of points that fit in the

town space while ensuring a sensible distance from each other. The town boundaries/biomes are created by a Voronoi segmentation [2] with Manhattan distance taken from each town position. This gives the segmentation a less triangular and more square appearance, that seems more natural with town shapes. Finally, the roads are plotted using an A* search [11], to the edge of each town and then from the edge of town to each buildings. The search avoids obstructions with a Chebyshev distance measure, to give square looking junctions/ring roads, and a Manhattan distance heuristic. Perlin noise [10] is used to provide the locations of other features, such as trees and lakes. The exact type of decoration (e.g. tree/lake/cactus) is dependent on the biome. This strategy is depicted in Figure 2, with a larger view of this layout from a top down perspective, without any visual assets. Town names are provided by a character based recurrent neural network trained on locations from ordinance survey data.

3.4 Motivation

LingoTowns option for a various number of games per town, allows for gamification elements to motivate players to contribute to annotation tasks which do not lend themselves as easily to gamification. The application of motivational game design mechanics in *LingoTowns* is, not solely to motivate annotation, which is the goal of the games it wraps, but primarily to direct a player's attention to provide a comprehensive annotation effort between tasks. We require that players complete a certain amount of each game as minimum to permit them to continue to the next level. Every game feeds back points into *LingoTowns*, with games that draw less attention delivering more points. In addition, special extra rewards are offered for games that players play less to encourage further completion.

3.5 Player Interaction

The player can interact with the map by clicking and dragging, using a two finger drag on a track pad, or using a screen touch and drag on a tablet or mobile device. When the preferred player location/town is no longer visible on screen, a pointer is displayed at the nearest edge of the screen showing the available town's location, level and completion percentage (shown in Figure 1). Each town has a floating dialog box above it describing the town, document and players completion so far. The player can click the buildings to visit the respective games for a town/document. Having played a game, the player is then returned to the world map.

3.6 Progression

In *LingoTowns* the progression is centralised, with *LingoTowns* deciding which document the player will see in each game. This arrangement allows us to address the cross-cutting interests of supporting learning between multiple tasks, maintaining player high engagement and gathering the annotations that are most needed. To achieve this we explore combining methods such as knowledge tracing with information theoretic methods derived from the aggregation. This provides a picture of player skills, while keeping the player engaged with the appropriate task difficulty, and ensuring a balance of placing the right player with the right task to optimise final annotation quality. For example, we can teach a player about nouns in *CafeClicker*, before moving them onto noun-phrases and mentions in *PhraseFarm*, and eventually coreference in *LingoTorium*. However, we can also incorporate more fine grained skills that look at the particular linguistic phenomena that occur on a grammatical level in the documents to tailor a learning experience that provides annotators with a combination of games and documents that delivers players an informative, yet sufficiently challenging journey as they play.

3.7 Towns and Games

LingoTowns is designed to support easy inclusion of third party games. These games can be played as part of *LingoTowns*, or independently. Currently, there are three sub games, namely: *CafeClicker*, a game for labelling part of speech tags; *PhraseFarm*, a game for labelling noun-phrases and *Lingotorium*, a game for labelling coreference. These games are an evolution of previous *WordClicker* [17], *TileAttack* [18] and *Wormingo* [13] respectively. The developments in the new games are the product of an in-depth user-study that focused on the player experience and user interaction. One of the core design challenges identified in our study common to all games, was the barrier to understanding created by the complexity of some of the tasks. In all games, special attention has been paid to the on-boarding process, with multiple opportunities for the player to discover more about the task, comprehension tests and more interactive tutorials.

3.7.1 CafeClicker. *CafeClicker*⁷ (shown in Figure 4) is a game for labelling part of speech tags, but could be expanded upon more broadly for annotation of lexical categories in general. The focus of the games changes following the user-study are largely related to clarity of feedback and cosmetic items. Game theme has changed slightly to better support the metaphor in way that more easily reflects the player’s actions. Instead of subtle changes to a cake, the player is now offered feedback based on a cake making machine.

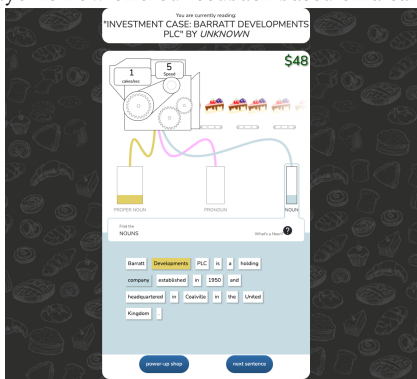


Fig. 4. *CafeClicker*

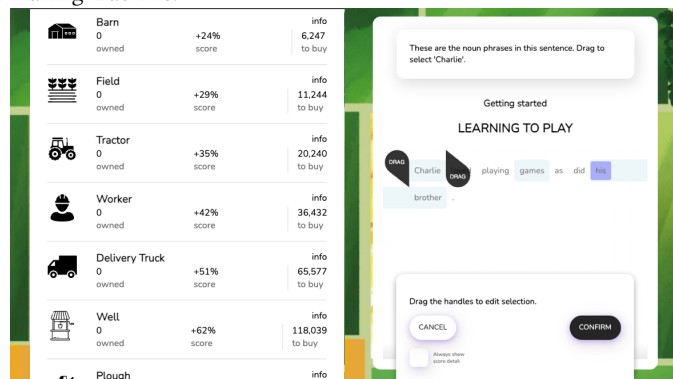


Fig. 5. *PhraseFarm*: Gameplay/Tutorial; Shop

3.7.2 PhraseFarm. *PhraseFarm*⁸ is a game used for noun-phrase annotation, but may be applied to any segmentation task. As players work through the sentences in the document, at the end of each sentence, we give detailed feedback on the annotations they provided in comparison with the annotations that we currently hold to ensure they continue to have a good understanding of the task as the complexity increases.

PhraseFarm now features an updated text segmentation interface that closely replicates the traditional cursor based text selection interface that users are used to seeing in a text editor to give a more familiar experience (shown in Figure 5).

Taking inspiration from recent research [17], like *CafeClicker*, *PhraseFarm* now features mechanics borrowed from the incremental game genre [1], such as the use of exponentially increasing points with an accompanying shop that allows players reinvest their points in farm themed items (Figure 5).

⁷<https://cafeclicker.com>

⁸<https://phrasefarm.org>

3.7.3 *Lingotorium*. *Lingotorium*⁹ serves as the final node of the *LingoTowns* pipeline, providing the annotations for coreference.

Lingotorium follows what we call the “motivation/annotation” paradigm [13], where players solve various word game-like puzzles (Figure 6) and annotation tasks back-to-back. This method aims to (1) enhance text comprehension as the puzzles require a certain level of understanding of the text in order to be solved (2) provide instant feedback as the system always knows the answers to the word-puzzles whereas the answer to a novel coreference may not be initially known (3) make the game more fun and engaging, especially players who enjoy word games.

Lingotorium gameplay is highly customisable on the motivational puzzles aspect. Players can adjust the difficulty, frequency of puzzle types or even turn puzzles off completely. To manage annotation difficulty, the game features a progression system [12] where players unlock more complicated tasks upon reaching a certain accuracy and number of annotations on the prior less complicated tasks (currently from discourse-old, to non-referring labelling and to plural antecedents tasks respectively [Figure 7])

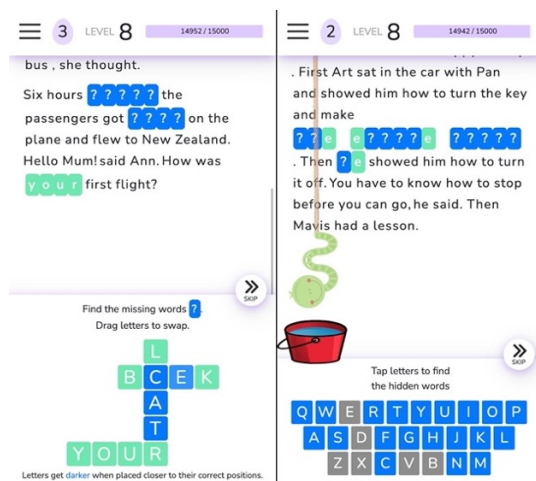


Fig. 6. *Lingotorium*: Motivational word puzzles

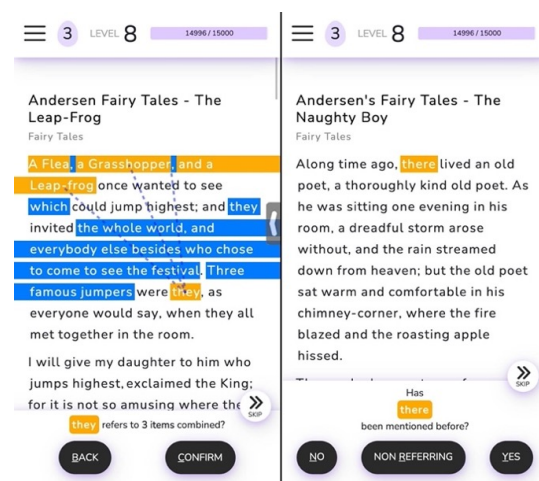


Fig. 7. *Lingotorium*: Plural antecedents and non-referring annotation tasks

4 CONCLUSION

In this work we have discussed our approach to organising multiple GWAPs to form a cooperative effort to annotate documents. *LingoTowns* serves as a base for exploring multiple concepts, including progression, motivation and gamification between multiple tasks and documents. *LingoTowns* already has three games, but we hope to collaborate with other GWAPs that require multi-task annotations or extend the existing natural language processing pipeline. In further work we will discuss the results of our approach on annotation and player enjoyment, and introduce new methods of progression and game-like motivation that build on the *LingoTowns* base. In the long term we hope to be able to release multiply annotated corpora from *LingoTowns* for supervised learning and study.

4.1 Acknowledgements

This research was supported by the DALI project, ERC Grant 695662.

⁹<https://lingotorium.com>

REFERENCES

- [1] Sultan A Alharthi, Olaa Alsaedi, Zachary O Toups, Theresa Jean Tanenbaum, and Jessica Hammer. 2018. Playing to wait: A taxonomy of idle games. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [2] Franz Aurenhammer. 1991. Voronoi diagrams—a survey of a fundamental geometric data structure. *ACM Computing Surveys (CSUR)* 23, 3 (1991), 345–405.
- [3] Richard A Bartle. 2004. *Designing virtual worlds*. New Riders.
- [4] Federico Bonetti and Sara Tonelli. 2020. A 3D Role-Playing Game for Abusive Language Annotation. In *Workshop on Games and Natural Language Processing*. 39–43.
- [5] Jon Chamberlain, Massimo Poesio, and Udo Kruschwitz. 2009. A new life for a dead parrot: Incentive structures in the Phrase Detectives game. In *Proceedings of the Webcentives Workshop at WWW*, Vol. 9.
- [6] Robert L Cook. 1986. Stochastic sampling in computer graphics. *ACM Transactions on Graphics (TOG)* 5, 1 (1986), 51–72.
- [7] Seth Cooper, Christoph Sebastian Deterding, and Theo Tsapakos. 2016. Player rating systems for balancing human computation games: testing the effect of bipartiteness. In *DIGRA/FDG'16-Proceedings of the First International Joint Conference of DIGRA AND FDG*. DIGRA Digital Games and Research Association.
- [8] Albert T Corbett and John R Anderson. 1994. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction* 4, 4 (1994), 253–278.
- [9] Vickie Curtis. 2015. Motivation to participate in an online citizen science game: A study of Foldit. *Science Communication* 37, 6 (2015), 723–746.
- [10] David S Ebert, F Kenton Musgrave, Darwyn Peachey, Ken Perlin, and Steven Worley. 2003. *Texturing & modeling: a procedural approach*. Morgan Kaufmann.
- [11] Peter E Hart, Nils J Nilsson, and Bertram Raphael. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics* 4, 2 (1968), 100–107.
- [12] Osman Doruk Kicikoglu, Richard Bartle, Jon Chamberlain, Silviu Paun, and Massimo Poesio. 2020. Aggregation Driven Progression System for GWAPs. In *Workshop on Games and Natural Language Processing*. 79–84.
- [13] Osman Doruk Kicikoglu, Richard Bartle, Jon Chamberlain, and Massimo Poesio. 2019. Wormingo: a 'true gamification' approach to anaphoric annotation. In *Proceedings of the 14th International Conference on the Foundations of Digital Games*. 1–7.
- [14] Raph Koster. 2013. *Theory of fun for game design*. " O'Reilly Media, Inc."
- [15] Mathieu Lafourcade. 2007. Making people play for Lexical Acquisition with the JeuxDeMots prototype. In *SNLP'07: 7th international symposium on natural language processing*. 7.
- [16] Mathieu Lafourcade, Alain Joubert, and Nathalie Le Brun. 2015. *Games with a Purpose (GWAPs)*. Wiley.
- [17] Christopher Madge, Richard Bartle, Jon Chamberlain, Udo Kruschwitz, and Massimo Poesio. 2019. Incremental game mechanics applied to text annotation. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*. 545–558.
- [18] Chris Madge, Jon Chamberlain, Udo Kruschwitz, and Massimo Poesio. 2017. Experiment-driven development of a GWAP for marking segments in text. In *Extended Abstracts Publication of the Annual Symposium on Computer-Human Interaction in Play*. 397–404.
- [19] Chris Madge, Juntao Yu, Jon Chamberlain, Udo Kruschwitz, Silviu Paun, and Massimo Poesio. 2019. Progression in a language annotation game with a purpose. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 7. 77–85.
- [20] Timothy Miller, Dmitriy Dligach, and Guergana Savova. 2012. Active learning for coreference resolution. In *BioNLP: Proceedings of the 2012 Workshop on Biomedical Natural Language Processing*. 73–81.
- [21] Alice Millour, Marianne Grace Araneta, Ivana Lazić Konjik, Annalisa Raffone, Yann-Alan Pilatte, and Karën Fort. 2019. Katana and Grand Guru: a Game of the Lost Words. In *9th Language & Technology Conference: Human Language Technologies as a Challenge for Computer Science and Linguistics (LTC'19)*.
- [22] Massimo Poesio, Jon Chamberlain, Udo Kruschwitz, Livio Robaldo, and Luca Ducceschi. 2013. Phrase detectives: Utilizing collective intelligence for internet-scale language resource creation. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 3, 1 (2013), 1–44.
- [23] M Jordan Raddick, Georgia Bracey, Pamela L Gay, Chris J Lintott, Carie Cardamone, Phil Murray, Kevin Schawinski, Alexander S Szalay, and Jan Vandenberg. 2013. Galaxy Zoo: Motivations of citizen scientists. *arXiv preprint arXiv:1303.6886* (2013).
- [24] Kathleen Tuite. 2014. GWAPs: Games with a problem.. In *FDG*.
- [25] Daniele Vannella, David Jurgens, Daniele Scarfini, Domenico Toscani, and Roberto Navigli. 2014. Validating and extending semantic knowledge bases using video games with a purpose. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 1294–1304.
- [26] Noortje Venhuizen, Kilian Evang, Valerio Basile, and Johan Bos. 2013. Gamification for word sense labeling. In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013)*.
- [27] Luis Von Ahn. 2006. Games with a purpose. *Computer* 39, 6 (2006), 92–94.
- [28] Luis von Ahn. 2013. Duolingo: learn a language for free while helping to translate the web. In *Proceedings of the international conference on Intelligent User Interfaces (IUI)*. ACM, 1–2. <https://doi.org/10.1145/2449396.2449398>
- [29] Luis Von Ahn and Laura Dabbish. 2008. Designing games with a purpose. *Commun. ACM* 51, 8 (2008), 58–67.

- [30] Michelle Yuan, Patrick Xia, Benjamin Van Durme, and Jordan Boyd-Graber. 2021. Adaptive Active Learning for Coreference Resolution. *arXiv preprint arXiv:2104.07611* (2021).