

“Hunt Takes Hare”: Theming Games Through Game-Word Vector Translation

Anonymous

Anonymous

ABSTRACT

A game’s theme is an important part of its design – it conveys narrative information, rhetorical messages, helps the player intuit strategies, aids in tutorialisation and more. Thematic elements of games are notoriously difficult for AI systems to understand and manipulate, however, and often rely on large amounts of hand-written interpretations and knowledge. In this paper we present a technique which connects *game embeddings*, a recent method for modelling game dynamics from log data, and *word embeddings*, which models semantic information about language. We explain two different approaches for using game embeddings in this way, and show evidence that game embeddings enhance the linguistic translations of game concepts from one theme to another, opening up exciting new possibilities for reasoning about the thematic elements of games in the future.

CCS CONCEPTS

• **Applied computing** → **Arts and humanities**; • **Computing methodologies** → **Artificial intelligence**.

KEYWORDS

procedural content generation, automated game design, computational creativity

ACM Reference Format:

Anonymous. 2024. “Hunt Takes Hare”: Theming Games Through Game-Word Vector Translation. In *Proceedings of the 19th International Conference on the Foundations of Digital Games (FDG 2024)*, May 21–24, 2024, Worcester, MA, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3649921.3659851>

1 INTRODUCTION

Almost a decade ago, Cook and Smith wrote that the ‘mechanics-first view on games is unnecessarily limiting’ and called for AI research into game design to consider experiential, aesthetic and rhetorical aspects of games in more depth [4]. Despite new work along these lines, building automated systems that can understand, change or create the theming and contextual meaning for a videogame is an underexplored area of game AI research.

Understanding a game’s theme requires the ability to relate game concepts to real-world ideas. This is useful for a wide variety of

applications, including the contextualisation of generated content, game design assistance and guidance, and tutorialisation. However, an abstract game carries no indication of how its components should be themed, and a surface-level translation of an existing theme that does not take into account how a game is played will likely lose coherence and impact.

Previous attempts to overcome this problem have tended to rely on pre-made databases of meaning, and user-defined graphs of concepts that relate to one another [14][15]. We cover some of this work in section §3, *Related Work*. While this approach has many benefits, it is very time-consuming, does not scale or generalise, and cannot be extended to procedurally generated concepts for which no prior knowledge exists.

In [12], Rabii et al present a technique which uses word embeddings to train a vector-space representation of game log data. They show that with no prior knowledge about the game other than the logs that it is possible to extract complex, high-level knowledge about the game’s structure, dynamics and strategy. This approach circumvents some of the issues mentioned above, in that it produces knowledge about a game simply through observations of it being played. However, this knowledge is completely abstract – it has no connection to the grounded ‘meaning’ of the game.

In this paper we repurpose Rabii’s work, and present our technique for using a combination of game log embeddings and word embeddings to translate a game’s theme across a semantic space. We describe our system setup and how we relate game embeddings to the real-world linguistic knowledge base represented by word embeddings. We also describe our approaches to translating thematic elements across the word embedding space to yield new themes for games. We share some preliminary results here, and note the challenges we have encountered so far, and our preliminary steps to improve upon the work. This represents a starting point for a new application of word embeddings to help relate game dynamics to real-world ideas.

The remainder of this paper is organised as follows: in section §2 we cover how word and game embeddings work; in §3 we discuss related work in computational creativity and automated game design; in §4 we cover the methodology of our approach to thematic translation; in §5 we present and evaluate initial results from the system; and in §6 we conclude our work and look to the future.

2 BACKGROUND

2.1 Word Embeddings

Word embeddings are representations of sequential data (usually language) in the form of n -dimensional vectors, used in natural language processing for tasks such as topic modelling and semantic distance measurement. The process is trained on a dataset of token sequences, and the positions of tokens in the resulting vector space aims to mimic their distributions in the original data. For natural

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FDG ’24, May 21–24, 2024, Worcester, Massachusetts, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0955-5/24/05...\$15.00

<https://doi.org/10.1145/3649921.3659851>

language, this means that words with a similar semantic role in the training data tend to appear close to each other along certain dimensions in the embedded space. Word embeddings are useful not only for measuring the distance between tokens, but also because mathematical operations can be performed on the vectors, such as addition, subtraction or averaging. This allows for calculations to be performed on words to compose or subtract meanings. Parrish gives the example of colours in [10]: subtracting the vector for the word *Blue* from the vector for the word *Sky* and adding the result to the vector for the word *Grass* yields a vector close to the word *Green*. The best-known algorithm for creating word embeddings is Word2Vec, which has been widely studied both within computer science as well as beyond in digital arts and other creative fields [9]. Word2Vec's simplicity and broad scope has made it a useful and accessible tool for artists and creative coders as well.

2.2 Game Embeddings

In [12] Rabii and Cook present an application of word embeddings to gameplay data. They take a set of formally-annotated chess match logs, and translate it to a domain-specific language that they designed for the purposes of training an embedding. They then apply the word2vec algorithm to this chess gameplay data, yielding a vector-space embedding for the game logs. In their paper they show that the resulting embedding reveals a range of interesting information about chess, from fundamental structural concepts through to subtle strategic insights understood by proficient players.

In their original proposal, the authors note several important traits of their system: the design of the domain-specific language, the size of the data, the source of data points (for example, expert players versus novices) and the exact phrasing of queries all have a significant impact on the types of conclusions that can be drawn from the data. Despite this, the work clearly shows that the embeddings capture important truths about the game simply by seeking semantic connections between logs of game events. This technique is potentially very scalable and widely-applicable as a result, since a game developer can easily gather event logs for their own game, and create a vector embedding from it.

Both game and word embeddings represent relationships between concepts through a multi-dimensional vector space, learned from datasets. As stated above, we can operate within these vector spaces to transform and translate concepts, through processes like addition or subtraction. However, vector spaces can also be related to one another, particularly when we can provide points which overlap or connect semantically. We can draw connections between these spaces through the use of regression techniques [13], which allow us to use known connections between spaces to infer new ones. In this paper, we present our efforts to draw a connection between a game embedding space (trained on chess data, as per Rabii and Cook) and a word embedding space (trained on English language text) to show how linguistic and ludic concepts can be linked together, and then transformed to find new themes and meaning in word embedding space.

3 RELATED WORK

3.1 Metaphor and Analogy

In this paper we seek to re-theme a game. We have a set of game elements, such as the pieces on a chessboard and the states and rules of the game, and we wish to rename these elements to align with a new theme. This has some similarity to the notion of metaphor or analogy in natural language processing research. A metaphor has two elements: a *tenor*, the root concept that is being conveyed, and a *vehicle*, a second concept that is used to represent the tenor in a new, metaphorical context. Metaphors usually connect concepts together through a specific connection point and then encourage the audience to extrapolate other, looser connections from there.

Metaphor and analogy are a widely studied topic in broader AI and natural language processing research, especially in computational creativity where researchers such as Veale have extensively studied how computers can extract metaphors from language [16]. Veale's approach identifies common structures in natural language and mining those associations to create a large knowledge base. This is supported by existing linguistic knowledge graphs such as WordNet [5]. In [6], Gero et al use the open-source knowledge graph GloVe to evaluate potential metaphors and grades different qualities of the metaphor based on their semantic similarity in the knowledge graph. Our approach also uses GloVe as a basis for its linguistic knowledge.

3.2 Theming Games

Theming or otherwise providing semantic meaning to game elements is primarily a topic covered by automated game design research, which Cook defines as 'the science and engineering of AI systems that model, participate in or support the game design process' [3]. Such systems engage with the topic of game theming for a variety of reasons, including wanting to impose a theme on an existing game, or interpreting a game to understand possible themes it may be expressing. Treanor's work on micro-rhetorics for the Game-o-Matic system is one such example – micro-rhetorics are user-defined components of meaning which can be combined with one another to compose higher-level rhetorical messages [15]. In his work on the Game-o-Matic, Treanor combines micro-rhetorics with thematic graphs that users provide (for example, 'cop arrests protester') to synthesise games which convey a theme.

GEMINI furthers work in this area, by providing bidirectional interpretation [14]. This allows the GEMINI system to not only create games based on thematic structures, but also to intuit possible thematic structures by examining an unthemed game, which is closer to the task we approach in this paper. This is done using a pre-compiled database of structures and their rhetorical interpretations, which formulates the task of theming a game to a logic programming task using answer set solvers. Notably, GEMINI operates on rules, not behaviour – 'mechanics', rather than 'dynamics', in the parlance of Hunnicke et al [7] – and its rhetorical databases are also defined in this way, describing the rhetorical meaning of game elements directly. Our approach considers the dynamic behaviour of game elements instead, by focusing on Rabii's game embeddings technique, and leverages large existing knowledge databases (in our case, word embedding models) that are not specific to any particular game.

Turn Black Rook C4 R7 C4 R0 Capture Knight	<i>Black moved a Rook from (4,7) to (4,0), capturing a Knight.</i>
Turn Black Pawn C5 R1 C5 R0 Promote Queen Check	<i>Black moved a Pawn from (5,1) to (5,0), promoting it to a Queen. The game is in check.</i>
Turn White King C4 R7 C2 R7 Castling Checkmate WhiteWin	<i>White moved its King from (4,7) to (2,7), with the single-use castling move. The game is in checkmate. White wins.</i>

Table 1: Excerpts of chess games in our description language (left) and their natural language equivalent (right)

Token	Words	Token	Words
White	white	Capture	capture
Black	black	Castling	castling
King	king	Promote	promote transform
Queen	queen	Checkmate	checkmate
Bishop	bishop	Check	check, control, prevent
Rook	rook	Stalemate	stalemate, deadlock
Knight	knight	WinWhite	victory
Pawn	pawn	WinBlack	defeat
		Draw	draw, tie, deadlock

Table 2: List of tokens and their associated words used to define Chess’s narrative theme

4 METHODOLOGY

Our goal is to build a system that can associate words in English to Chess game concepts, by relating English word embeddings to game embeddings built from Chess play data. If *token* is a game token, we note G_{token} its associated game vector. If *word* is a word in our dataset’s English vocabulary, we note W_{word} its associated word vector. If w_1, w_2, \dots, w_n are words in our dataset’s English vocabulary, we note W_{w_1, w_2, \dots, w_n} the average word vector $\sum_{i=1}^n \frac{w_i}{n}$. The relationship vector between two words w_1 and w_2 is denoted $W_{w_1 \rightarrow w_2} = W_{w_2} - W_{w_1}$. This is a vector which can be ‘added’ to the vector W_{w_1} to yield the vector representing W_{w_2} .

4.1 Training Embeddings

Our source for word vectors are the public domain GloVe pre-trained embeddings, trained on the Wikipedia 2014 and the Gigaword 5 corpus [11]. Its vocabulary contains 400,000 words, each represented by a 50-dimensional vector.

To create our game embeddings, we reproduced Rabii and Cook’s [12] data processing pipeline. We first downloaded a month of ranked matches data from the online Chess playing platform Lichess, and converted each logged game in the custom description language described by the authors. Samples of the converted data with their natural language equivalent are provided in Table 1.

The description language contains a total of 34 different tokens that represent different concepts used to describe a game of Chess. Tokens can represent players (e.g. White, Black), Chess pieces (Queen, Pawn), playing moves (Capture, Castling), game states

(Checkmate, WinBlack) and even the eight rows and columns of the board (e.g. R0, R1, R7, C0, C1, C7).

We then used our corpus of Chess data to train a Word2Vec model using the same settings as Rabii and Cook. In the end, each of our 34 Chess tokens has a 5-dimensional vector representation.

4.2 Themes and Rethemings

We define a *theming* as a function that associates game tokens built from our Chess description language to word vectors in our dataset.

We consider Chess to have an already established theming, which we denote as the function *theme*. Our definition of Chess’ current theme can be found in Table 2. It does not contain every token in the original description language, only those that we could associate to a specific word vector. For example, since we associate the game token King to the word "king", we have:

$$theme(King) = W_{king}$$

However, *theme* cannot take as input the tokens that represent the rows and columns of the board (R0, . . . , C7), as we couldn’t map them to a specific word in English. Some tokens can be associated to a word, but that word might have other meanings than the one we want to use (e.g. Draw does not refer to the act of drawing a picture, but a situation where no player wins). In that case, we associate the token to at most three words w_1, \dots, w_3 , and use the average of their word vectors W_{w_1, w_2, w_3} . The set of tokens that are valid inputs for *theme* is called \mathcal{T} .

Our goal is to *retheme* Chess, i.e. create other themings with alternative token-to-words associations. We define *retheme* as a function that, given a token and a relationship vector $start \rightarrow end$, returns a word vector:

$$retheme_{start \rightarrow end}(token) = W_{word}$$

The relationship vector guides the semantic translation of the input token into the final output word. For example, if we built a function to shift the masculine words in Chess’ theming to their feminine equivalents, we might have the following equations:

$$theme(King) = W_{king}$$

$$retheme_{masculine \rightarrow feminine}(King) = W_{queen}$$

In the following sections, we describe three different models for building the *retheme* function, one using only word vectors, and two using a combination of word and game vectors. In the following results section, we show and discuss results from all three approaches.

4.3 Word Vectors Only

Our first retheming model is built using the well-known Word2vec’s analogy formula, adapted to the general case:

$$retheme_{start \rightarrow finish}(token) = theme(token) + W_{start \rightarrow finish}$$

For the above example of translating across the theme vector of masculine to feminine, we would therefore calculate the retheming as:

$$W_{queen} = W_{king} + W_{masculine \rightarrow feminine}$$

This function only uses word vectors. We chose it as a baseline to compare with techniques that leverage the information stored in game vectors.

4.4 Using Both Game and Word Vectors

Game embeddings contain expert knowledge about the specific game they were extracted from, while word embeddings represent the meaning of words in the broad context of their dataset. For example, some expert players assign a numeric valuation to each chess piece as a way to quantify how important they are for winning a game. If we compute the cosine similarity between G_{King} , G_{Queen} , ..., G_{Pawn} and the vector $G_{Checkmate}$, we can extract an ordering of chess pieces which aligns with the ordering of importance as perceived by chess masters. When doing the same basic operation with word vectors instead, we found no sign of correlation with the expert valuation of chess pieces. Table 3 shows these comparisons. This suggests that game embedding data carries important information about the structure of Chess as a *game* which the words alone do not. Therefore, we hypothesize that a retheming using both word and game vectors can provide us with more relevant results than ones purely based on word vectors.

We want our model to be able to translate relationships between game tokens (such as the relative value ordering of chess pieces) into relationships between words (e.g. the medieval hierarchy of Chess piece names). We are interested in models of the form:

$$\text{retheme}_{start \rightarrow finish}(\text{token}) = f(G_{\text{token}}) + W_{start \rightarrow finish}$$

where f is a linear function that associates a game vector to a word vector. The choice of linear functions is motivated by the fact that semantic relationships between words and game concepts are captured in relationship vectors $G_{1 \rightarrow 2}$ or $W_{1 \rightarrow 2}$. [12] [2] A linear transformation $f : X \rightarrow \alpha X + \beta$ preserves relationship vectors up to a scaling factor α .

$$f(G_{1 \rightarrow 2}) = f(G_2) - f(G_1) = \alpha W_2 - \alpha W_1 = \alpha W_{1 \rightarrow 2}$$

Recall that \mathcal{T} is the set of tokens which are valid inputs for the theming function. For each game token $T \in \mathcal{T}$, we train a linear regression model f_T on a subset $\mathcal{T}_{N,T}$ of N game tokens randomly chosen in \mathcal{T} and their corresponding value given by *theme*. The purpose of f_T is to convert the gameplay relationships embedded in the game vector G_T to semantic relationships, encoded in a word vector. Unlike the method presented in Section 4.3, we don't want the prediction to be influenced by the currently given value of $\text{theme}(G_T)$ so we make sure that the token T is absent from $\mathcal{T}_{N,T}$.

$$\mathcal{T}_{N,T} \subseteq \mathcal{T} \quad T \notin \mathcal{T}_{N,T} \quad |\mathcal{T}_{N,T}| = N$$

$$(\alpha_T, \beta_T) = \underset{(\alpha, \beta)}{\text{Argmin}} \sum_{t \in \mathcal{T}_{N,T}} \|\alpha G_t + \beta - \text{theme}(t)\|^2$$

$$f_T(G_T) = \alpha_T G_T + \beta_T$$

Once the conversion from the game to word embeddings space is done, we add the relationship vector that guides the retheming:

$$\text{retheme}_{start \rightarrow finish}(T) = f_T(G_T) + W_{start \rightarrow finish}$$

4.5 Choosing the guiding vector

The *retheme* function requires a relationship word vector of the form $W_{start \rightarrow finish} = W_{finish} - W_{start}$, that we call the guiding vector. It is best to think of it as representing a semantic transformation from the starting word *start* to the target word *finish*. In our experiments, this transformation is applied to the narrative context of Chess, to shift it from its current theme to a different target theme. We experimented with two types of guiding vector.

4.5.1 Guiding based on a specific example. If we have a token in our Chess description language that we want to retheme to a specific word (eg. map the token *King* to the word *lion* instead of the word *king*), we can provide the vector representing that exact transformation: $W_{king \rightarrow lion}$. All the other rethemed word vectors will be translated by that same vector. We can think of this translation as being anchored around the given example, with the other translations being secondary and hopefully following the same overall trajectory.

4.5.2 Guiding based on a target semantic field. Instead of providing a specific retheming for a token (*king* \rightarrow *lion*), we can choose a broader semantic field, represented by a list of words $B = b_1, b_2, \dots, b_n$. Let $A = a_1, a_2, \dots, a_m$ be the words used in the current theming of Chess in Table 2. We can create a vector representing the shift from one lexical field to another by first computing the average vector of each list – respectively $W_B = \frac{1}{n} \sum_{i=1}^n W_{b_i}$ and $W_A = \frac{1}{m} \sum_{i=1}^m W_{a_i}$. The guiding vector $W_{A \rightarrow B} = W_B - W_A$ represents the transformation of shifting from the starting theme to the one specified by the list of words given in input.

5 RESULTS

The models described in Section §4 were built with the goal of retheming the tokens in our Chess description language. We present our results under different modalities in Tables 4 and 5. Each model is tasked with associating all the game tokens in Table 2 to a word embedding. We omit the exact values of each 50-dimensional vector here and instead give the top three closest words featured in our English dictionary, according to cosine distance.

We explore two types of retheming, each with the goal of shifting the original medieval warfare theming of Chess to a theme about wildlife. Table 4 contains the models output when prompted with the specific transformation $W_{king \rightarrow lion}$ (see 4.5.1). Table 5 contains results obtained when prompting the models with the target semantic field represented by the following words: *lion, elephant, zebra, eating, becoming, extinction, control, win, loss*. Our explorations showed that ideally, this list should not only contain words that refer to animals, as words about the basic concepts of a strategy game (eg. *win* and *loss*) are important to define the semantic field of a Chess game.

Applying the procedure outlined in 4.5.2 gave us a guiding vector $W_{A \rightarrow B}$ that is similar to the vector $W_{face \rightarrow wild}$, as the closest words to the average of each semantic field were "*face*" for Chess, and "*wild*" for the list of words related to wildlife.

For each modality, we compare the results of three models: a baseline retheming using only word vectors, followed by two rethemings using both word and game vectors combined, with the sampling parameter N set to 5 and 10, respectively.

Token	Queen	Rook	King	Bishop	Knight	Pawn
Expert valuation (Evans, 1958)	10.0	5.0	4.0	3.75	3.5	1.0
Similarity to token "Checkmate" (game vectors)	10.0	9.96	9.41	5.99	4.93	1.0
Similarity to token " <i>checkmate</i> " (word vectors)	2.91	7.36	3.49	1.0	5.64	10.0
Similarity to token " <i>win</i> " (word vectors)	9.61	1.0	10.0	4.42	9.84	5.58

Table 3: All values are linearly normalised to be between 1 and 10.

Tokens	1. Word vectors only	2. Word+Game (N=10)	3. Word+Game (N=5)
Black	blue black pink	foxes grizzly hyena	blue white green
White	blue black green	lion grizzly bear	black blue pink
Bishop	saltire borough primate	hyena heelers furry	grayish striped silvery
King	lion dragon elephant	swan queen princess	continentwide free-range bizonal
Knight	lion sable grizzly	striped underside blue	suffragan ripon archbishop
Pawn	pawn man-sized henhouse	lion sable hyena	black blue green
Queen	lion mermaid dragon	whelp bullseye passant	black blue green
Rook	rook dippin maned	bear small black	lion lady queen
Check	mosquito trap sniffing	lion boar dragon	lion swan dragon
Checkmate	whelp t-1000 airborne	snow white climate	preyed genovese villainous
Draw	fringe u.s.-controlled netting	staving staved dampener	pawn gajar dispater
Stalemate	staving u.s.-controlled continentwide	capering side-chain polyelectrolyte	pawn gumdrop bejeweled
WinBlack	underdogs wolves relegation	tied tie bottom	stormy impasse mbeki
WinWhite	victory 1-0 2-1	underdogs relegation rout	decisive defeat rout
Capture	capturing trap capture	lion grey sable	sable lion bloods
Castling	bow-tie ant-like tzitzit	lion rainbow blue	black blue white
Promote	empowerment sustainable promotes	ringlets hydrophobic quartile	metahumans angoulême 54-kg
Average R ²	N/A	0.61 ± 0.06	1.0 ± 0.0

Table 4: Rethemings of game tokens using guiding vector *King* → *Lion*

Tokens	1. Word vectors only	2. Word+Game (N=10)	3. Word+Game (N=5)
Black	wild black squirrel	wild boar ant	mormon latter-day darwin
White	wild green white	wild domesticated ant	indle great-grandson ahrts
Bishop	latter-day apostle bishop	wild foxes hunters	knight grandson nephew
King	king famous ancestor	cat lion butterflies	spread areas weeks
Knight	hunter badger wolf	boar thornberrys pheasant	capture capturing expedition
Pawn	pawn boar stinkhorn	wild hare native	dragon knight adventures
Queen	queen enchanted famous	wild hunting hunters	spearpoint ftsz n-channel
Rook	rook groening voboril	hunt wild famous	capture capturing elude
Check	wild breeding cattle	bee gees narnia	famous queen wild
Checkmate	eriogonum brettanomyces aphid	deadlock stalemate impasse	cattle wild sheep
Draw	wild ducks 2-2	wild disastrous logging	black white gray
Stalemate	spawning stalemate ftaa	bcg spiritism self-sustaining	darwin augustine lds
WinBlack	wild defeat 4-0	victory wild 4-0	rook ajaw vålerengen
WinWhite	victory wild 4-0	wild defeat tigers	mccurry mcclellan fratto
Capture	hunters wild hunt	ant wild asteraceae	mittelfranken r2000 k21
Castling	castling pseudotooth mouflon	wild ghost hunter	nelson clive muir
Promote	breeding ecotourism transgenic	tackie amed jan.-may	stalemate deadlock impasse
Average R ²	N/A	0.6 ± 0.03	1.0 ± 0.0

Table 5: Rethemings of game tokens using a target semantic field related to "wildlife"

The maximum value for N is the size of the original Chess theming in Table 2 minus one, i.e. 16. Our experiments using values of N close to 16 yielded poor results, with little variation between each output. Choosing lower values of N provides more diversity in the output word vectors, but increases the probability for each linear regression model to over-fit on the training set, which is evidenced by their average R^2 of 1 in the lower case where $N = 5$. Over-fitting reduces the model’s ability to generalize, leading to more words that are either too close to the original theme (such as *knight* or *queen*) or seem too far from the target theme (eg. *mormon*, *mittel – franken*, *mccurry*). We observed empirically that good trade-off seem to lie in the middle, with $N = 10$. This trade-off may be different for other games and for differently-sized or -distributed token sets, which future work will investigate.

Comparing the baseline rethemings (Column 1) and the ones with $N = 10$ (Column 2), we observe that models that don’t use game vectors often provides words that are close, if not exactly the same as the starting theme (*king* for King, *pawn* for Pawn, *apostle* for Bishop) while the ones that do use them reuse few words from the original Chess theme.

For both modalities, the guiding vector’s target words ("*lion*" and "*wild*" respectively) are often the closest word to the output of each model. This behavior was already observed in previous work on Word2Vec analogies [8] [1]. As a result of this work, a common practice in analogy tasks is to simply discard target words from the results and use the second closest instead. We follow this principle in Step 2 of our retheming algorithm (5.1).

5.1 Comparing Retheming of Chess Pieces

We compare our $N = 10$ model and the baseline model on the task of generating a set of chess pieces themed around wildlife. We use the results generated in Table 5, columns 1 and 2. For each token related to a chess piece, we follow this algorithm:

- (1) Compute the three closest words to the model’s output.
- (2) If the guiding vector’s target word is in the list, discard it.
- (3) Among the remaining words, pick the first noun of the list.

The result of that process is presented in Table 6. We note that the baseline model outputs words that are quite similar from the starting theme (Column 1), while the model using game vectors seems to correctly feature a "wildlife" theme (Column 2).

In Column 2, we observe that the medieval social hierarchy of Chess has been replaced by a food chain: weaker pieces are associated to prey animals while stronger pieces are associated to predators. Interpreting the mapping, we remark that the King is linked to a *cat*, an animal that is both weak and cherished, which bears some similarity to the King’s role in chess.

The Rook is associated with *hunt*, while the Bishop is associated with *foxes*. The Queen’s moveset allows it to move either like a Bishop or a Rook, and its associated word *hunters* could be considered a composition of both the notion of a group hunt, represented by the Rook, and foxes which are solitary hunting animals, represented by the Bishop. Semantic blends such as this are necessarily fuzzy, but we find these examples inspirational and engaging.

We consider that our model’s ability to provide a retheming whose semantic relationships seems to match, at least in an interpretive manner, with gameplay properties is a strong sign of its

Token	Word vectors only	Themed (Word+Game)
King	king	cat
Queen	queen	hunters
Bishop	apostle	foxes
Rook	rook	hunt
Knight	hunter	boar
Pawn	pawn	hare

Table 6: Outputs of our retheming models, when tasked with generating a Chess set that fits a "wildlife" theme

potential. The baseline model doesn’t exhibit the same properties, suggesting that exploiting the information captured in game embeddings is valuable to explore the link between gameplay and narrative theming.

6 DISCUSSION AND CONCLUSIONS

In this paper we presented preliminary results on the use of both word and game embeddings to model and translate thematic elements of games. We motivated our work by building on prior research into game embeddings for chess, and word embeddings for natural language, and showed how the two can be combined together in different ways to alter the theme of a game such as chess. We provided evidence that our technique works better than simply using word embeddings alone to retheme words related to a game, and we propose that this is because the game embeddings capture essential qualities of various game design components and this knowledge supports a richer translation.

We believe this approach has many promising and interesting features. For one, it captures the *dynamics* of a game, rather than its static elements: game embeddings reveal, for example, the relative strength of chess pieces, even though this knowledge is not described anywhere in the game’s rules. Most automated game design research focuses on static mechanical elements of games, and struggles to bridge to dynamic or emergent properties. This could be an exciting new way to approach the understanding, extraction and use of dynamic game elements.

Embeddings are also scalable and comparatively easy to both create and understand. Adding logging to a game to record important events and create datasets, such as the one we used from chess matches, is a simple exercise no different to the debugging or playtesting that game developers already do. We have already worked with one independent game developer to add logging of this type to their game, and the process was clean and straightforward. Word2Vec arithmetic is a widely-used and taught technique among digital art communities and creative coders. This can be furthered through better tools build specifically to make training and exploring embedding spaces more amenable. This makes it a promising technique for widespread adoption in the games industry.

We are planning to experiment further with this technique on more traditional digital games, to assess how it performs on different kinds of game scenario and different kinds of game logs. At the time of writing, we are working with an independent game developer to gather data from one of their games, from both expert and novice players. We also hope to experiment with more ambitious uses for the technique as well, such as the ability to create

names for game elements that have been procedurally generated at runtime. We believe there are many exciting applications for this basic technique waiting to be discovered.

ACKNOWLEDGMENTS

We would like to thank the reviewers for their helpful feedback and reflections on the paper. While we weren't able to include all of their suggestions, this is an ongoing work and we will continue to develop the ideas and approaches here in future. This work was supported by the IGGI CDT, and the Royal Academy of Engineering.

REFERENCES

- [1] Carl Allen and Timothy Hospedales. 2019. Analogies explained: Towards understanding word embeddings. In *International Conference on Machine Learning*. PMLR, 223–231.
- [2] Carl Allen and Timothy M. Hospedales. 2019. Analogies Explained: Towards Understanding Word Embeddings. *CoRR* abs/1901.09813 (2019). arXiv:1901.09813 <http://arxiv.org/abs/1901.09813>
- [3] Mike Cook. 2022. Puck: A Slow and Personal Automated Game Designer. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*.
- [4] Michael Cook and Gillian Smith. 2015. Formalizing Non-Formalism: Breaking the Rules of Automated Game Design. In *Proceedings of the Foundations of Digital Games Conference*.
- [5] Christiane Fellbaum (Ed.). 1998. *WordNet: An Electronic Lexical Database*. MIT Press, Cambridge, MA.
- [6] Katy Ilonka Gero and Lydia B. Chilton. 2019. Metaphoria: An Algorithmic Companion for Metaphor Creation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery.
- [7] Robin Hunicke, Marc Leblanc, and Robert Zubek. 2004. MDA: A Formal Approach to Game Design and Game Research. In *Workshops of the AAAI Conference on Artificial Intelligence*.
- [8] Omer Levy and Yoav Goldberg. 2014. Linguistic regularities in sparse and explicit word representations. In *Proceedings of the eighteenth conference on computational natural language learning*. 171–180.
- [9] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *arXiv:1301.3781 [cs]* (Sept. 2013). <http://arxiv.org/abs/1301.3781> arXiv: 1301.3781.
- [10] Allison Parrish. 2018. Understanding Word Vectors. github.com/aparrish.
- [11] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*. 1532–1543. <http://www.aclweb.org/anthology/D14-1162>
- [12] Younès Rabii and Mike Cook. 2021. Revealing Game Dynamics via Word Embeddings of Gameplay Data. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*.
- [13] PEDRO L. RODRIGUEZ, ARTHUR SPIRLING, and BRANDON M. STEWART. 2023. Embedding Regression: Models for Context-Specific Description and Inference. *American Political Science Review* 117, 4 (2023), 1255–1274. <https://doi.org/10.1017/S0003055422001228>
- [14] Adam Summerville, Chris Martens, Ben Samuel, Joseph Osborn, Noah Wardrip-Fruin, and Michael Mateas. 2018. Gemini: Bidirectional Generation and Analysis of Games via ASP. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*.
- [15] Mike Treanor, Bobby Schweizer, Ian Bogost, and Michael Mateas. 2012. The micro-rhetorics of Game-o-Matic. In *Proceedings of the International Conference on the Foundations of Digital Games*.
- [16] Tony Veale, Ekaterina Shutova, Beata Beigman Klebanov, Graeme Hirst, and Carlo Strapparava. 2018. Metaphor: A Computational Perspective. *Computational Linguistics* 44, 1 (2018).