

# Sounds Wilde

Phonetically extended embeddings for author-stylized poetry generation

Alexey Tikhonov

Yandex

Karl-Liebknecht strasse 1, Berlin

altsoph@gmail.com

Ivan P. Yamshchikov

Max Planck Institute

for Mathematics in the Sciences

Inselstrasse 22, Leipzig

ivan@yamshchikov.info

## Abstract

This paper addresses author-stylized text generation. Using a version of a language model with extended phonetic and semantic embeddings for poetry generation we show that phonetics has comparable contribution to the overall model performance as the information on the target author. Phonetic information is shown to be important for English and Russian language. Humans tend to attribute machine generated texts to the target author.

## 1 Introduction

Generative models for natural languages and for poetry specifically are discussed for more than fifty years (Wheatley, 1965). Lamb et al. (2017) provides a detailed overview of generative poetry techniques. This particular paper addresses the issue of author stylized poetry (Tikhonov and Yamshchikov, 2018a) and shows the importance of phonetics for the author-stylized poetry generation.

The structure of the poem can vary across different languages starting with highly specific and distinct structures of Chinese poems (Zhang et al., 2017) and finishing with poems where formal structure hardly exists, e.g. American poetry of the twentieth century (say, the lyrics of Charles Bukowski) or so-called *white poems* in Russian. The structure and standards of poems can depend on various factors primarily phonetic in their nature. In the broadest sense, rhymes in a classical western sonnet, a structure of a Persian ruba'i, a sequence of tones in a Chinese quatrain or a structure within rap bars could be expressed as a set of phonetic rules based on a certain understanding of expressiveness and euphony shared across a given culture or, sometimes, an artistic community. Some cultures and styles also have particular semantic limitations or 'standards', for

example, 'centrality' of certain topics in classical Japanese poetry, see (Maynard, 1994). We do not make attempts to address high-level semantic structure, however one can add some kind of pseudo-semantic rules to the model discussed further, say via some mechanism in line with (Ghazvininejad et al., 2016) or (Yi et al., 2017). The importance of phonetics in poetical texts was broadly discussed among Russian futuristic poets, see (Kruchenykh, 1923). Several Russian linguistic circles and art groups (particularly OPOJAZ) in the first quarter of 20th century were actively discussing the concept of the *abstruse* language, see (Shklovsky, 1919), stressing also that the form of a poem, and especially its acoustic structure, is a number one priority for the future of literature. In their recent paper (Blasi et al., 2016) have challenged the broadly accepted idea that sound and meaning are not interdependent: unrelated languages very often use (or avoid) the same sounds for specific referents. In (He et al., 2016) and (Ghannay et al., 2016) it was shown that acoustic word embeddings improve algorithm performance on a number of NLP tasks. In line with these ideas, one of the key features of the model that we discuss below is its concatenated embedding that contains information on the phonetics of every word preprocessed by a bi-directional Long Short-Term Memory (LSTM) network alongside with its vectorized semantic representation.

In (Tikhonov and Yamshchikov, 2018a) a model for generation of texts resembling the writing style of a particular author within the training data set was proposed. In this paper we quantify the stylistic similarity of the generated texts and show the importance of extension of the word embeddings with phonetic information for the overall performance of the model. The proposed model might also be applicable to prose, but diverse phonetic structure of the poetry discussed above makes it

better suited for the purposes of this paper. Also, since one would like to incorporate human judgement of the generated text and measure how well a human reader can attribute generated text to the target author, poetry seems preferable to prose for its stylistic expressiveness.

The contribution of this paper is three-fold: (1) we propose an LSTM with extended phonetic and semantic embeddings and quantify the quality of the obtained stylized poems both subjectively through a survey and objectively with BLEU metrics; (2) we show that phonetic information plays key role in a author-stylized poetry generation (3) we demonstrate that the proposed approach works in a multilingual setting, providing examples in English and in Russian.

## 2 Related work

In (Lipton et al., 2015), (Kidson et al., 2016), (Lebret et al., 2016), (Radford et al., 2017), (Tang et al., 2016), (Hu et al., 2017) a number of RNN-based generative or generative adversarial models for controlled text generation are developed. These papers took *content* and *semantics* of the output into consideration, yet did not work with the style of the generated texts. In (Li et al., 2016) the authors focused on the speaker consistency in a dialogue. In (Sutskever et al., 2011) and in (Graves, 2013) it is demonstrated that a character-based recurrent neural network with gated connections or LSTM networks respectively can generate texts that resemble news or Wikipedia articles. Chinese classical poetry due to its diverse and deeply studied structure is addressed in (He et al., 2012), (Yi et al., 2017), (Yan, 2016), (Yan et al., 2016), or (Zhang et al., 2017). In (Ghazvininejad et al., 2016) an algorithm generates a poem in line with a user-defined topic in (Potash et al., 2015) stylized rap lyrics are generated with LSTM trained on a rap poetry corpus.

There is a diverse understanding of literary style that lately became obvious due to the growing attention to the problems of automated style transfer. For a brief overview of the state-of-the-art style transfer problem see (Tikhonov and Yamshchikov, 2018b). Style is sometimes regarded as a sentiment of a text (see (Shen et al., 2017) or (Li et al., 2018)), it’s politeness (Sennrich et al., 2016) or *style of the time* (see (Hughes et al., 2012)). In (Fu et al., 2017) authors generalize these ideas measuring the success of a particular style aspect with

a specifically trained classifier. In (Guu et al., 2017) it is shown that an existent human-written source used to control the saliency of the output can significantly improve the quality of the resulting texts. Generative models on the other hand often do not have such input and have to generate stylized texts from scratch, like in (Ficler and Goldberg, 2017).

## 3 Model

We use an LSTM-based language model that predicts the  $w_{n+1}$  word based on  $w_1, \dots, w_n$  previous inputs and some other parameters of the modeled sequence. A schematic picture of the model is shown in Figure 1, document information projections obtained as a matrix multiplication of document embedding on a projection matrix the dimensionality of which differs according to the target dimensionality of a projection are highlighted with blue arrows. An LSTM with 1152-dimensional input and 512-dimensional state.

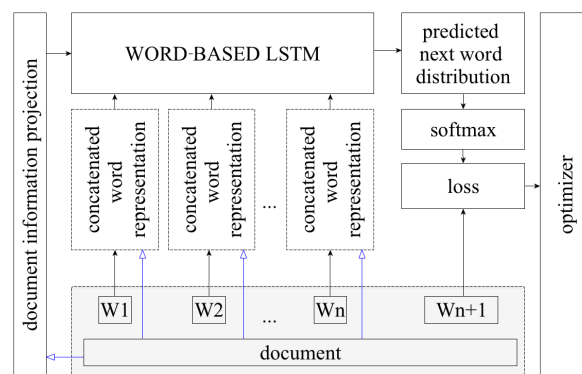


Figure 1: The scheme of the language model used. Document information projections are highlighted with blue arrows. The projections on a state space of the corresponding dimension is achieved with simple matrix multiplication of document embeddings.

Figure 2 shows a concatenated word representation of the model. The representation includes a 512-dimensional projection of a concatenated author and document embeddings at every step and two 128-dimensional vectors corresponding to final states of two bidirectional LSTMs. The first LSTM works with a char-representation of the word and the second one uses phonemes of the International Phonetic Alphabet<sup>1</sup>, employing an heuristics to transcribe words into phonemes. A somewhat similar idea, but with convolutional neural networks rather than with LSTMs, was pro-

<sup>1</sup><http://www.internationalphoneticalphabet.org>

	N. of documents	Size of vocab.	N. of authors	Size
English	110000	165000	19000	150 Mb
Russian	330000	400000	1700	140 Mb

Table 1: Parameters of the training datasets.

	N. of words		N. of words
Shakespeare	10 218	Pushkin	226 001
Carroll	19 632	Esenin	73 070
Marley	22 504	Letov	29 766
MUSE	7 031	Zemfira	23 099

Table 2: Number of words in the training datasets for human-peer tested lyrics.

posed in (Jozefowicz et al., 2016).

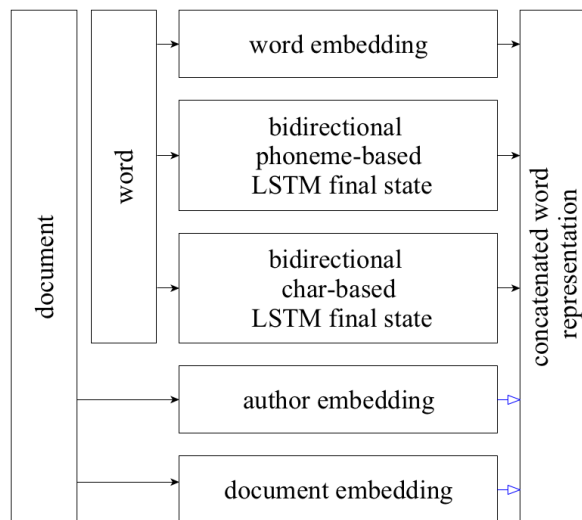


Figure 2: Concatenated word representation.

## 4 Datasets

Two proprietary datasets of English and Russian poetry were used for training. All punctuation was deleted, every character was transferred to a lower case. No other preprocessing was made. The datasets sizes can be found in Table 2. This allowed to have approximately 330 000 verses in train dataset and another 10 000 verses forming a test dataset for Russian poetry. For English poetry train data consisted of 360 000 verses with approximately 40 000 verses forming the test data.

The model was trained for English (William Shakespeare, Edgar Allan Poe, Lewis Carroll, Oscar Wilde and Bob Marley as well as lyrics of the American band Nirvana and UK band Muse) and Russian (Alexander Pushkin, Sergey Esenin,

Joseph Brodsky, Egor Letov and Zemfira Ramazanova).

The model produces results of comparable quality for both languages, so in order to make this paper shorter, we further address generative poems in English only and provide the experimental results for Russian in the Appendix. We want to emphasize that we do not see any excessive difficulties in implementation of the proposed model for other languages for which one can form a training corpus and provide a phonetically transcribed vocabulary.

Table 3 shows some generated stylized poetry examples. The model captures syntactic characteristics of the author (note the double negation in the first and the last line of Neuro-marley) alongside with the vocabulary ('burden', 'darkness', 'fears' could be subjectively associated with gothic lyrics of Poe, whereas 'sunshine', 'fun', 'fighting every rule' could be associated with positive yet rebellious reggae music). Author-specific vocabulary can technically imply specific phonetics that characterizes a given author, however this implication is not self evident and generally speaking does not have to hold. As we demonstrate later, phonetics does, indeed, contribute to the author stylization significantly.

## 5 Experiments

In (Tikhonov and Yamshchikov, 2018a) the detailed description of the experiments is provided alongside with a new metric for automated stylization quality estimation — *sample cross entropy*. In this submission we specifically address the results that deal with the phonetics of the generated texts.

### 5.1 BLEU

BLEU is a metric estimating the correspondence between a machine's output and that of a human and is usually mentioned in the context of machine translation. We suggest to adopt it for the task of stylized text generation in the following way: a random starting line is sampled out of the human-written poems and initializes the generation. Generative model 'finishes' the poem generating the ending lines of the quatrain. Then one calculates BLEU between three actual lines that finished the human-written quatrain starting with a given first line and lines generated by the model when initialized with the same human-written line.

Table 4 shows BLEU calculated on the validation dataset for the plain vanilla LSTM, LSTM

Neuro-Poe	Neuro-Marley
her beautiful eyes were bright	don t you know you ain t no fool
this day is a burden of tears	you r gonna make some fun
the darkness of the night	but she s fighting every rule
our dreams of hope and fears	ain t no sunshine when she s gone

Table 3: Examples of the generated stylized quatrains. The punctuation is omitted since it was omitted in the training dataset.

with author information support but without bidirectional LSTMs for phonemes and characters included in the embeddings and the full model. The uniform random and weighted random give baselines to compare the model to.

Model $G(A_i)$	BLEU
Uniform Random	0.35%
Weighted Random	24.7%
Vanilla LSTM	29.0%
Author LSTM	29.3% (+1% to vanilla LSTM)
Full model	29.5% (+1.7% to vanilla LSTM)

Table 4: BLEU for uniform and weighted random random sampling, vanilla LSTM, LSTM with author embeddings but without phonetics, and for the full model. Phonetics is estimated to be almost as important for the task of stylization as the information on the target author.

Table 4 shows that extended phonetic embeddings play significant role in the overall quality of the generated stylized output. It is important to mention that phonetics in an implicit characteristic of an author and the training dataset (in line with the definition of style in (Tikhonov and Yamshchikov, 2018b)), humans do not have qualitative insights into phonetic of Wilde or Cobain, yet the information on it turns out to be important for the style attribution.

## 5.2 Survey data

Two quatrains from William Shakespeare, Lewis Carroll, Bob Marley and MUSE band were sampled. They were accompanied by two quatrains generated by the model conditioned on those four authors respectively. One hundred and forty fluent English-speakers were asked to read all 16 quatrains in randomized order and choose one option out of five offered for each quatrain, i.e. the author of this verse is William Shakespeare, Lewis Carroll, Bob Marley, MUSE or an Artificial Neural Network. The summary of the obtained results is shown in Table 5. Analogous results but for Russian language could be seen in Appendix in Table 8 alongside with more detailed description of the

methodology. It is important to note that the generated pieces for tests were human-filtered for mistakes, such as demonstrated in Table 6, whereas the automated metrics mentioned above were estimated on the whole sample of generated texts without any human-filtering.

Looking at Table 5 one can see the model has achieved good results in author stylization. Indeed the participants recognized Shakespeare more than 46% of the times (almost 2.5 times more often than compared with a random choice) and did slightly worse in their recognition of Bob Marley (40% of cases) and MUSE (39% of cases, still 2 times higher than a random choice). This shows that the human-written quatrains were, indeed, recognizable and the participants were fluent enough in the target language to attribute given texts to the correct author. At the same time, people were 'tricked' into believing that the text generated by the model was actually written by a target author in 37% of cases for Neuro-Shakespeare, 47% for Neuro-Marley, and 34% for Neuro-MUSE, respectively. Somehow, Lewis Carroll turned out to be less recognizable and was recognized in the survey only in 20% of cases (corresponds to a purely random guess). The subjective underperformance of the model on this author can therefore be explained with the difficulty experienced by the participants in determining his authorship.

Combining the results in Table 4 with the results of the survey shown in Table 5 one could conclude that phonetic structure of lyrics has impact on the correct author attribution of the stylized content. This impact is usually not acknowledged by a human reader but is highlighted with the proposed experiment.

## 6 Conclusion

In this paper we have addressed a problem of author-stylized text generation. It has been shown that the extending word embeddings with phonetic information has a comparable impact on the BLEU

Model $G(A_i)$ or author	Shakespeare	Carroll	Marley	MUSE	LSTM
Neuro-Shakespeare	<b>0.37*</b>	0.04	0.05	0.14	<b>0.3**</b>
Shakespeare	<b>0.46*</b>	0.05	0.04	0.07	<b>0.3**</b>
Neuro-Carroll	0.02	0.07	<b>0.26**</b>	0.18	<b>0.41*</b>
Carroll	0.05	<b>0.2**</b>	0.14	0.11	<b>0.32*</b>
Neuro-Marley	0.02	0.01	<b>0.47*</b>	0.2	<b>0.29**</b>
Marley	0.15	0.05	<b>0.4*</b>	0.1	<b>0.24**</b>
Neuro-MUSE	0.09	0	0.12	<b>0.34**</b>	<b>0.39*</b>
MUSE	0.03	0.05	<b>0.28**</b>	<b>0.39*</b>	0.2

Table 5: Results of a survey with 140 respondents. Shares of each out of 5 different answers given by people when reading an excerpt of a poetic text by the author stated in the first column. The two biggest values in each row are marked with \* and \*\* and a bold typeface.

of the generative model as the information on the authors of the text. It was also shown that, when faced with an author with a recognizable style (an author who is recognized approximately two times more frequently than at random), humans mistakenly recognize the output of the proposed generative model for the target author as often as they correctly attribute original texts to the author in question. The experiments were carried out in English and in Russian and there are no obvious obstacles for the application of the same approach to other languages.

## References

- E Blasi, Damián, Søren Wichmann, Harald Hammarström, F. Stadler, Peter, and H. Christiansen, Morten. 2016. Sound – meaning association biases evidenced across thousands of languages. *Proceedings of the National Academy of Sciences*, 113(39):10818 – 10823.
- Jessica Fidler and Yoav Goldberg. 2017. Controlling linguistic style aspects in neural language generation. In *Proceedings of the Workshop on Stylistic Variation*, pages 94 – 104.
- Zhenxin Fu, Xiaoye Tan, Nanyun Peng, Dongyan Zhao, and Rui Yan. 2017. Style transfer in text: Exploration and evaluation. In *arXiv preprint:1711.06861*.
- Sahar Ghannay, Yannick Estève, Nathalie Camelin, and Deléglise Paul. 2016. Evaluation of acoustic word embeddings. In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP 2016*, pages 62–66.
- Marjan Ghazvininejad, Xing Shi, Yejin Choi, and Kevin Knight. 2016. Generating topical poetry. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1183–1191. Association for Computational Linguistics.
- Alex Graves. 2013. Generating sequences with recurrent neural networks. In *arXiv preprint:1308.0850*.
- Kelvin Guu, Tatsunori B. Hashimoto, Yonatan Oren, and Percy Liang. 2017. Generating sentences by editing prototypes. In *arXiv preprint:1709.08878*.
- Jing He, Ming Zhou, and Long Jiang. 2012. Generating chinese classical poems with statistical machine translation models. In *AAAI*.
- Wanjia He, Weiran Wang, and Karen Livescu. 2016. Multi-view recurrent neural acoustic word embeddings. In *arXiv preprint:1611.04496*.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. 2017. Toward controlled generation of text. In *International Conference on Machine Learning*, pages 1587–1596.
- James M. Hughes, Nicholas J. Foti, David C. Krakauer, and Daniel N. Rockmore. 2012. Quantitative patterns of stylistic influence in the evolution of literature. *Proceedings of the National Academy of Sciences*, 109(20):7682–7686.
- Rafal Jozefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. 2016. Exploring the limits of language modeling. In *arXiv preprint:1602.02410*.
- Chloe Kiddon, Luke Zettlemoyer, and Yejin Choi. 2016. Globally coherent text generation with neural checklist models. *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 329–339.
- Aleksei Kruchenykh. 1923. *Phonetics of theater*. M.:41, Moscow.
- Carolyn Lamb, G. Brown, Daniel, and L. Clarke, Charles. 2017. A taxonomy of generative poetry techniques. *Journal of Mathematics and the Arts*, 11(3):159–179.
- Remi Lebret, David Grangier, and Michael Auli. 2016. Neural text generation from structured data with application to the biography domain. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1203–1213.

- Jiwei Li, Michel Galley, Chris Brockett, Georgios P. Spithourakis, Jianfeng Gao, and William B. Dolan. 2016. A persona-based neural conversation model. *CoRR*, abs/1603.06155.
- Juncen Li, Robin Jia, He He, and Percy Liang. 2018. Delete, retrieve, generate: A simple approach to sentiment and style transfer. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 1, pages 1865–1874.
- Zachary C. Lipton, Sharad Vikram, and Julian McAuley. 2015. Capturing meaning in product reviews with character-level generative text models. In *arXiv preprint:1511.03683*.
- Senko K. Maynard. 1994. The centrality of thematic relations in japanese text. *Functions of language*, 1(2):229–260.
- Peter Potash, Alexey Romanov, and Anna Rumshisky. 2015. Ghostwriter: Using an lstm for automatic rap lyric generation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1919–1924. Association for Computational Linguistics.
- Alec Radford, Rafal Jozefowicz, and Ilya Sutskever. 2017. Learning to generate reviews and discovering sentiment. In *arXiv preprint:1704.01444*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Controlling politeness in neural machine translation via side constraints. *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 35–40.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. *31st Conference on Neural Information Processing Systems*, pages 6833–6844.
- Boris Shklovsky. 1919. *Poetics: on the theory of poetic language*. 18 State typography, Petrograd.
- Ilya Sutskever, James Martens, and Geoffrey Hinton. 2011. Generating text with recurrent neural networks. In *Proceedings of the 28th International Conference on Machine Learning*, pages 1017–1024.
- Jian Tang, Yifan Yang, Sam Carton, Ming Zhang, and Qiaozhu Mei. 2016. Context-aware natural language generation with recurrent neural networks. In *arXiv preprint:1611.09900*.
- Alexey Tikhonov and Ivan P. Yamshchikov. 2018a. Guess who? multilingual approach for the automated generation of author-stylized poetry. In *arXiv preprint:1807.07147*.
- Alexey Tikhonov and Ivan P. Yamshchikov. 2018b. What is wrong with style transfer for texts? In *arXiv preprint:1808.04365*.
- Jon Wheatley. 1965. The computer as poet. *Journal of Mathematics and the Arts*, 72(1):105.
- Rui Yan. 2016. i, poet: Automatic poetry composition through recurrent neural networks with iterative polishing schema. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, pages 2238–2244.
- Rui Yan, Cheng-Te Li, Xiaohua Hu, and Ming Zhang. 2016. Chinese couplet generation with neural network structures. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 2347 – 2357.
- Xiaoyuan Yi, Ruoyu Li, and Maosong Sun. 2017. Generating chinese classical poems with rnn encoder-decoder. In *Chinese Computational Linguistics and Natural Language Processing Based on Naturally Annotated Big Data*, pages 211–223.
- Jiyuan Zhang, Yang Feng, Dong Wang, Yang Wang, Andrew Abel, Shiyue Zhang, and Andi Zhang. 2017. Flexible and creative chinese poetry generation using neural memory. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 1364–1373.

## A Examples of output

Table 6 lists some illustrative mistakes of the model both for English and for Russian language. Reading the raw output we could see several types of recurring characteristic errors that are typical for LSTM-based text generation. They can be broadly classified into several different types:

- If the target author is underrepresented in the training dataset, model tends to make more mistakes, mostly, syntactic ones;
- Since generation is done in a word-by-word manner, model can deviate significantly when sampling a low-frequency word;
- Pronouns tend to cluster together, possibly due to the problem of anaphoras in the training dataset;
- The line can end abruptly, this problem also seems to occur more frequently for the authors that are underrepresented in the training dataset.

Table 7 lists some subjectively cherry-picked especially successful examples of the system outputs both for English and for Russian language. Since text is generated line by line and verses are obtained through random rhyme or rhythm filtering several types of serendipitous events occur. They can be broadly classified into four different types:

<b>Problem</b>	<b>English</b>	<b>Russian</b>
<b>broken syntax</b>	<i>Neuro – MUSE :</i> every step inside on our faces that i would stop my self going crazy	<i>Neuro – Zemfira :</i> ты слышишь я слышу шаги мои в душу дрожи
<b>'rare' word brakes the line</b>	<i>Neuro – Shakespeare :</i> o ho de profundis she says i am on her	<i>Neuro – Letov :</i> иду гляжу в окно гляжу мне вслед на небо
<b>pronouns 'entangle'</b>	<i>Neuro – Shakespeare :</i> thou here shalt be and thine who will have to my grave	<i>Neuro – Zemfira :</i> тебе ли ты ль за
<b>sentences don't end</b>	<i>Neuro – Muse :</i> at night i lay waiting for a	<i>Neuro – Brodsky :</i> двух четырех десять лет за углом
<b>nonsense</b>	<i>Neuro – Shakespeare :</i> do many a fair honour best of make or lose	<i>Neuro – Lenov :</i> о о о и о

Table 6: Examples of several recurring types of mistakes that occur within generated lyrics.

- Wording of the verse that fits into the style of the target author;
- Pseudo-plot that is perceived by the reader due to a coincidental cross-reference between two lines;
- Pseudo-metaphor that is perceived by the reader due to a coincidental cross-reference between two lines;
- Sentiment and emotional ambience that correspond to the subjective perception of the target author.

## B Survey design

The surveys were designed identically for English and Russian languages. We have recruited the respondents via social media, the only prerequisite was fluency in the target language. Respondents were asked to determine an authorship for 16 different 4-line verses. The verses for human-written

text were chosen randomly out of the data for the given author. The generated verses were obtained through line-by-line automated rhyme and rhythm heuristic filtering. Since LSTMs are not perfect in text generation and tend to have clear problems illustrated in Table 6 we additionally filtered generative texts leaving the verses that do not contain obvious mistakes described above. Each of the 16 questions consisted of a text (in lower case with a stripped-off punctuation) and a multiple choice options listing five authors, namely, four human authors and an artificial neural network. Respondents were informed that they are to distinguish human- and machine-written texts. The correct answers were not shown to the respondents. Table 5 shows the results of the survey for English texts and Table 8 for Russian ones. Higher values in every row correspond to the options that were more popular among the respondents, when they were presented with the text written by the author listed in the first column of the table.

Serendipity	English	Russian
<b>peculiar wording</b>	<i>Neuro – Shakespeare :</i> a sense i may not comprehend of whom i had not to defend	<i>Neuro – Pushkin :</i> во славу вакха или тьмы мы гордо пиروвали
<b>apophenic plot</b>	<i>Neuro – Marley :</i> oh lord i know how long i d burn take it and push it it s your turn	<i>Neuro – Esenin :</i> ты под солнцем стоишь и в порфире как в шелку беззаботно горишь
<b>apophenic metaphor</b>	<i>Neuro – Carroll :</i> your laugh is bright with eyes that gleam that might have seen a sudden dream	<i>Neuro – Zemfira :</i> ветер в голове с красной тенью шепчется
<b>peculiar sentiment</b>	<i>Neuro – Muse :</i> i am the man of this universe i remember i still am a curse	<i>Neuro – Letov :</i> только в ушах отражается даль только белая смерть превращается в ад

Table 7: Cherry-picked examples of generated lyrics after either rhyme or rhythm filtering illustrating typical serendipities.

Model $G(A_i)$ or author	Pushkin	Esenin	Letov	Zemfira	LSTM
<b>Neuro-Pushkin</b>	<b>0.31**</b>	0.22	0.02	0.0	<b>0.44*</b>
<b>Pushkin</b>	<b>0.62*</b>	0.11	0.03	0.01	<b>0.23**</b>
<b>Neuro-Esenin</b>	0.02	<b>0.61*</b>	0.08	0.0	<b>0.29**</b>
<b>Esenin</b>	0.06	<b>0.56*</b>	0.07	0.02	<b>0.29**</b>
<b>Neuro-Letov</b>	0.0	0.02	<b>0.40**</b>	0.08	<b>0.51*</b>
<b>Letov</b>	0.0	0.01	<b>0.61*</b>	0.02	<b>0.35**</b>
<b>Neuro-Zemfira</b>	0.0	0.06	0.13	<b>0.4**</b>	<b>0.41*</b>
<b>Zemfira</b>	0.0	0.02	0.08	<b>0.58*</b>	<b>0.31**</b>

Table 8: Results of a survey with 178 respondents. Shares of each out of 5 different answers given by people when reading an exempt of a poetic text by the author stated in the first column. The two biggest values in each row are marked with \* and \*\* and a bold typeface.