

Effect of Language and Error Models on Efficiency of Finite-State Spell-Checking and Correction

Tommi A Pirinen

University of Helsinki
Department of Modern Languages
FI-00014 Univ. of Helsinki, PO box 24
tommi.pirinen@helsinki.fi

Sam Hardwick

University of Helsinki
Department of Modern Languages
FI-00014 Univ. of Helsinki, PO box 24
sam.hardwick@helsinki.fi

Abstract

We inspect the viability of finite-state spell-checking and contextless correction of non-word errors in three languages with a large degree of morphological variety. Overlooking previous work, we conduct large-scale tests involving three languages — English, Finnish and Greenlandic — and a variety of error models and algorithms, including proposed improvements of our own. Special reference is made to on-line three-way composition of the input, the error model and the language model. Tests are run on real-world text acquired from freely available sources. We show that the finite-state approaches discussed are sufficiently fast for high-quality correction, even for Greenlandic which, due to its morphological complexity, is a difficult task for non-finite-state approaches.

1 Introduction

In most implementations of spell-checking, efficiency is a limiting factor for selecting or discarding spell-checking solutions. In the case of finite-state spell-checking it is known that finite-state language models can efficiently encode dictionaries of natural languages (Beesley and Karttunen, 2003), even for polysynthetic languages. Most contemporary spell-checking and correction systems are still based on programmatic solutions (e.g. hunspell¹, and its *spell relatives), or at most specialised algorithms for implementing error-tolerant traversal of the finite-state dictionaries (Oflazer, 1996; Huldén, 2009a). There have also been few fully

finite-state implementations that both detect and correct errors (Schulz and Mihov, 2002; Pirinen and Lindén, 2010). In this paper we further evaluate the use of finite-state dictionaries with two-tape finite-state automata as a mechanism for correcting misspellings, and optimisations to the finite-state error models, intending to demonstrate that purely finite-state algorithms can be made sufficiently efficient.

To evaluate the general usability and efficiency of finite-state spell-checking we test a number of possible implementations of such a system with three languages of typologically different morphological features² and *reference implementations for contemporary spell-checking applications*: English as a morphologically more isolating language with essentially a word-list approach to spell-checking; Finnish, whose computational complexity has been just beyond the edge of being too hard to implement nicely in e.g. hunspell (Pitkänen, 2006); and Greenlandic, a polysynthetic language which is implemented as a finite-state system using Xerox’s original finite-state morphology formalism (Beesley and Karttunen, 2003). As a general purpose finite-state library we use HFST³, which also contains our spell-

¹<http://hunspell.sf.net>

²We will not go into details regarding the morphological features of these languages. We thank the anonymous reviewer for guiding us to make a rough comparison using a piece of translated text. We observe from the translations of the *Universal Declaration of Human Rights* (with pre-amble included) as follows: the number of word-like tokens for English is 1,746, for Finnish 1,275 and for Greenlandic 1,063. The count of the 15 most frequent tokens are for English 120—28, for Finnish 85—10 and for Greenlandic 38—7.

The average word length is 5.0 characters for English, 7.8 for Finnish and 14.9 for Greenlandic. For the complexity of computational models refer to Table 2 in this article.

³<http://hfst.sf.net>

checking code.

As neither Finnish nor Greenlandic have been successfully implemented in the hunspell formalism, we mainly use them to evaluate how the complexity of a language model affects the efficiency of finite-state spell-checking. For a full-scale survey on the state-of-the-art non-finite-state spell-checking, refer to Mitton (2009).

The efficiency results are contrasted with the existing research on finite-state spell-checking in Hassan et al. (2008) and the theoretical results on finite-state error-models in Mitankin (2005). Our contribution primarily comprises the addition of morphologically complex languages with actual cyclic dictionary automata (i.e. infinite dictionaries formed by compounding and recurring derivation) and more complex structure in general, compared to those of English and Arabic. Our goal is to demonstrate that finite-state spelling is tractable for these complex languages, to document their implications for performance and to present an algorithm for the task. We also point out that previous approaches have neglected to simultaneously constrain the error model and the dictionary with each other in on-line composition, which affords a significant speed benefit compared to generating the two component compositions.

The rest of the paper is organised as follows. In Section 2 we discuss the spell-checking task, current non-finite-state spell-checkers and previously used finite-state methods for spell-checking and correction and propose some possible speed optimisations for the error models. We also investigate algorithmic limitations of finite-state approaches and ways to remedy them. In Section 3 we present the language models, error models and the testing corpora. In Section 4 we present the comparisons of speed and quality with combinations of different language and error models and corpora for spell-checking. In Section 5 we summarise our findings and results, and outline future goals.

2 Methods

A finite-state spell-checker is typically (Pirinen and Lindén, 2010) composed of at least two finite-state automata; one for the dictionary of the language, or the *language model*, which contains valid strings of

the language, and one automaton to map misspelt words into correct strings, or the *error model*. Both the language model and the error model are usually (Pirinen and Lindén, 2010) weighted finite-state automata, where the weights represent the probabilities are of a word being correctly spelled in the language model and of specific misspellings, respectively. We evaluate here the effect of both the language and error model automata's structure and complexity on the efficiency of the finite-state spelling task.⁴

2.1 Language Models

The most basic language model for a spell-checking dictionary is a list of correctly spelled word forms. One of the easiest ways of creating such a spell-checker is to collect the word forms from a reasonably large corpus of (mostly) correctly spelt texts. Additionally we can count the frequency of words and use that as the likelihood, $P(w) = \frac{c(w)}{\sum_{w \in \mathcal{D}} c(w)}$ where $c(w)$ is the count of the word w and \mathcal{D} is the set of corpus word forms. For morphologically more isolating languages such as English, this is often a sufficient approach (Norvig, 2010), and we use it to create a dictionary for our English spell-checker as well. As a non-finite-state reference point we use hunspell.

For agglutinative languages like Finnish, for which the word-list approach is likely to miss a much greater number of words, one of the most common approaches is to use right-linear grammars, possibly combined with finite-state rule languages to implement morphophonological alterations (Koskenniemi, 1983). This approach also applies to the newest available free / open source and full-fledged finite-state Finnish morphological dictionary we found (Pirinen, 2011). This language model features productive derivations, compounding and rudimentary probabilistic models. We take, as a reference non-finite state language model for Finnish, Voikko's implementation in Malaga, which is currently used as a spell-checking component in open source software. It is implemented in a

⁴The methods introduced in this research as well as all materials are free/libre open source. Please see our svn repository <https://hfst.svn.sf.net/svnroot/trunk/fsmnlp-2012-spellers/> for detailed implementation and scripts to reproduce all the results.

left-associative grammar formalism, which is a potentially less efficient system with more expressive power. It's similar to finite-state formulations in terms of linguistic coverage.

For polysynthetic languages it will be obvious that the coverage of any word-list-based approach will be even lower. Furthermore, most simple extensions to it such as affix stripping (as in hunspell) are not adequate for describing word forms. To our knowledge, the only approaches that have been widely used for spell-checking and morphological analysis of Greenlandic have been based on traditional finite-state solutions, such as the Xerox formalisms. In our case we have obtained a freely available finite-state morphology implementation from the Internet⁵. For further details we refer to the authors' website <http://oqaaserpassualeriffik.org/>.

2.2 Error Models

The ubiquitous formula for modeling typing errors since computer-assisted spelling correction began has been the edit distance metric sometimes attributed to Levenshtein (1966) and/or Damerau (1964). It maps four typical slips of the fingers on a keyboard to events in the fuzzy matching of misspelt word forms to correct ones, that is, the deletion of a character (i.e. failing to press a key), addition of a character (i.e. hitting an extra key accidentally), changing a character (i.e. hitting the wrong key) and transposing adjacent characters (i.e. hitting two keys in the wrong order).

When modeling edit distance as a finite-state automaton, a relatively simple two-tape automaton is sufficient to implement the algorithm (Hassan et al., 2008). The automaton will consist of one arc for each type of error, and additionally one state for each transposition pair. This means that the trivial nondeterministic finite-state automaton implementing the algorithm is of space complexity $S(V, E, \Sigma) = O(|\Sigma|^2|V| + |\Sigma|^2|E|)$, where Σ is the alphabet of language, V is the set vertices in automaton and E is the set of edges in automaton. This edit distance formulation is roughly feature equivalent to hunspell's TRY mechanism.

⁵<https://victorio.uit.no/langtech/trunk/st/kal>

To further fine-tune this finite-state formulation of the edit distance algorithm, it is possible to attach a probability to each of the error events as a weight in a weighted finite-state automaton, corresponding to the likelihood of an error, or a confusion factor. This can be used to implement features like keyboard adjacency or an OCR confusion factor to the error correction model. This will not modify the structure of the finite-state error models or the search space—which is why we did not test their effects in this article—but introduction of non-homogenous weights to the resulting finite-state network may have an effect on search time. This addition is equivalent to hunspell's KEY mechanism.

For English language spelling correction there is also an additional type of error model to deal with competence-related misspellings—as opposed to models that mainly deal with mistypings—implemented in the form of phonemic folding and unfolding. This type of error is very specific to certain types of English text and is not in the scope of this experiment. This is the PHON part of the hunspell's correction mechanism.

After fine-tuning the error models to reimplement hunspell's feature set, we propose variations of this edit distance scheme to optimise the speed of error correction with little or no negative effect to the quality of the correction suggestions. The time requirement of the algorithm is determined by the size of the search space, i.e. the complexity of the resulting network when the error model is applied to the misspelt string and intersected with the dictionary⁶.

To optimise the application of edit distance by limiting the search space, many traditional spell checkers will not attempt to correct the very first letter of the word form. We investigated whether this decision is a particularly effective way to limit the search space, but it does not appear to significantly differ from restricting edits at any other position in the input.

Dividing the states of a dictionary automaton into

⁶For non-finite-state solutions, the search space is simply the number of possible strings given the error corrections made in the algorithm. For finite-state systems the amount of generated strings with cyclic language and error models is infinite, so complexity calculation are theoretically slightly more complex, however for basic edit distance implementations used in this article the search space complexities are always the same and the amount of suggestions generated finite

classes corresponding to the minimum number of input symbols consumed by that state, we found that the average ambiguity in a particular class is somewhat higher for the first input symbols, but then stabilises quickly at a lower level. This was accomplished by performing the following state-categorisation procedure:

1. The start state is assigned to class 0, and all other states are assigned to a candidate pool.
2. All states to which there is an (input) epsilon transition from the start state are assigned to class 0 and removed from the candidate pool.
3. This is repeated for each state in class 0 until no more states are added to class 0. This completes class 0 as the set of states in which the automaton can be before consuming any input.
4. For each state in class 0, states in the candidate pool to which there is a non-epsilon transition are assigned to class 1 and removed from the candidate pool.
5. Class 1 is epsilon-completed as in (2-3).
6. After the completion of class n , class $n + 1$ is constructed. This continues until the candidate pool is empty, which will happen as long as there are no unreachable states.

With this categorisation, we tallied the total number of arcs from states in each class and divided the total by the number of states in the class. This is intended as an approximate measure of the ambiguity present at a particular point in the input. Some results are summarized in Table 1.

Class	Transitions	States	Average
0	156	3	52
1	1,015	109	9.3
2	6,439	1,029	6.3
3	22,436	5,780	3.9
4	38,899	12,785	3.0
5	44,973	15,481	2.9
6	47,808	17,014	2.8
7	47,495	18,866	2.5
8	39,835	17,000	2.3
9	36,786	14,304	2.6
10	45,092	14,633	3.1
11	66,598	22,007	3.0
12	86,206	30,017	2.9

Table 1: State classification by minimum input consumed for the Finnish dictionary

Further, the size of a dictionary automaton that is restricted to have a particular symbol in a particular

position does not apparently depend on the choice of position. This result was acquired by intersecting eg. the automaton $\epsilon . +$ with the dictionary to restrict the first position to have the symbol ϵ , the automaton $. \epsilon . +$ to restrict the second position, and so on. The transducers acquired by this intersection vary in size of the language, number of states and number of transitions, but without any trend depending on the position of the restriction. This is in line with the rather obvious finding that the size of the restricted dictionary in terms of number of strings is similarly position-agnostic.

Presumably, the rationale is a belief that errors predominately occur at other positions in the input. As far as we know, the complete justification for this belief remains to be made with a high-quality, hand-checked error corpus.

On the error model side this optimisation has been justified by findings where between 1.5 % and 15 % of spelling errors happen in the first character of the word, depending on the text type (Bhagat, 2007); the 1.5 % from a small corpus of academic texts (Yan-nakoudakis and Fawthrop, 1983) and 15 % from dictated corpora (Kukich, 1992). We also performed a rudimentary classification of the errors in the small error corpus of 333 entries from Pirinen et al. (2012), and found errors at the first position in 1.2 % of the entries. Furthermore, we noticed that when evenly splitting the word forms in three parts, 15 % of the errors are in the first third of the word form, while second has 47 % and third 38 %, which would be in favor of discarding initial errors⁷.

A second form of optimisation that is used by many traditional spell-checking systems is to apply a lower order edit distance separately before trying higher order ones. This is based on the assumption that the vast majority of spelling errors will be of lower order. In the original account of edit distance for spell-checking, 80 % of the spelling errors were found to be correctable with distance 1 (Pollock and Zamora, 1984).

The third form of optimisation that we test is omitting redundant corrections in error models of higher order than one. Without such an optimisa-

⁷By crude classification we mean that all errors were forced to one of the three classes at weight of one, e.g. a series of three consecutive instances of the same letters was counted as a deletion at the first position.

tion, higher order error models will permit adding and deleting the same character in succession at any position, which is obviously futile work for error correction. Performing the optimisation makes the error model larger but reduces the search space, and does not affect the quality of results.

2.3 Algorithms

The obvious baseline algorithm for the task of finding which strings can be altered by the error model in such a way that the alteration is present in the language model is generating all the possible alterations and checking which ones are present in the language model. This was done in Hassan et al. (2008) by first calculating the composition of the input string with the error model and then composing the result with the language model.

If we simplify the error model to one in which only substitutions occur, it can already be seen that this method is quite sensitive to input length and alphabet size. The composition explores each combination of edit sites in the input string. If any number of edits up to d can be made at positions in an input string of length n , there are

$$\sum_{i=1}^d \binom{n}{i}$$

ways to choose the edit site, and each site is subject to a choice of $|\Sigma| - 1$ edits (the entire alphabet except for the actual input). This expression has no closed form, but as d grows to n , the number of choices has the form 2^n , so the altogether complexity is exponential in input length and linear in alphabet size (quadratic if transpositions are considered).

In practice (when d is small relative to n) it is useful to observe that an increase of 1 in distance results in an additional term to the aforementioned sum, the ratio of which to the previously greatest term is

$$\frac{n!/(d! \cdot (n-d!))}{n!/((d-1)! \cdot (n-d+1)!)} = \frac{n-d+1}{d}$$

indicating that when d is small, increases in it produce an exponential increase in complexity. For an English 26-letter lowercase alphabet, edit distance 2 and the 8-letter word “spelling”, 700 strings are stored in a transducer. With transpositions, deletions, insertions and edit weights this grows to

100,215 different outputs. We have implemented this algorithm for our results by generating the edited strings by lookup, and performing another lookup with the language model on these strings.

Plainly, it would be desirable to improve on this. The intuition behind our improvement is that when editing an input string, say “spelling”, it is a wasted effort to explore the remainder after generating a prefix that is not present in the lexicon. For example, after changing the first character to “z” and not editing the second character, we have the prefix “zp-”, which does not occur in our English lexicon. So the remaining possibilities - performing any edits on the remaining 7-character word - can be ignored.

This is accomplished with a three-way composition in which the input, the error model and the language model simultaneously constrain each other to produce the legal correction set. This algorithm is presented in some detail in Lindén et al. (2012). A more advanced and general algorithm is due to Al-lauzen and Mohri (2009).

3 Material

For language models we have acquired suitable free-to-use dictionaries, readily obtainable on the Internet.

We made our own implementations of the algorithms to create and modify finite-state error models. Our source repository contains a Python script for generating error models and an extensive Makefile for exercising it in various permutations.

To test the effect of correctness of the source text to the speed of the spell-checker we have retrieved one of largest freely available open source text materials from the Internet, i.e. Wikipedia. The Wikipedia text is an appropriate real-world material as it is a large body of text authored by many individuals, and may be expected to contain a wide variety of spelling errors. For material with more errors, we have used a simple script to introduce (further, arbitrary) errors at a uniform probability of 1/33 per character; using this method we can also obtain a corpus of errors with correct corrections along them. Finally we have used a text corpus from a language different than the one being spelled to ensure that the majority of words are not in the vocabulary and (al-

most always) not correctable by standard error models.

The Wikipedia corpora were sourced from wikimedia.org. For exact references, see our previously mentioned repository. From the dumps we extracted the contents of the articles and picked the first 100,000 word tokens for evaluation.

In Table 2 we summarize the sizes of automata in terms of structural elements. On the first row, we give the size of the alphabet needed to represent the entire dictionary. Next we give the sizes of automata as nodes and arcs of the finite-state automaton encoding the dictionary. Finally we give the size of the automaton as serialised on the hard disk. While this is not the same amount of memory as its loaded data structures, it gives some indication of memory usage of the program while running the automaton in question. As can be clearly seen from the table, the morphologically less isolating languages do fairly consistently have larger automata in every sense.

Automaton	En	Fi	KI
Σ set size	43	117	133
Dictionary FSM nodes	49,778	286,719	628,177
Dictionary FSM arcs	86,523	783,461	11,596,911
Dictionary FSM on disk	2.3 MiB	43 MiB	290 MiB

Table 2: The sizes of dictionaries as automata

In Table 3 we give the same figures for the sizes of error models we’ve generated. The Σ size row here shows the number of symbols left when we have removed the symbols that are usually not considered to be a part of a spell-checking mechanism, such as all punctuation that does not occur word-internally and white-space characters⁸. Note that sizes of error models can be directly computed from their parameters; i.e., the distance, the Σ set size and the optimisation, so this table is provided for reference only.

4 Evaluation

We ran various combinations of language and error models on the corpora described in section 3. We give tabular results of the speed of the system and the effect of the error model on recall. The latter

⁸The method described here does not handle run-on words or extraneous spaces, as they introduce lot of programmatic complexity which we believe is irrelevant to the results of this experiment.

Automaton	En	Fi	KI
Σ set size	28	60	64
Edit distance 1 nodes	652	3,308	3,784
Edit distance 1 arcs	2,081	10,209	11,657
Edit distance 2 nodes	1,303	6,615	7,567
Edit distance 2 arcs	4,136	20,360	23,252
No firsts ed 1 nodes	652	3,308	3,784
No firsts ed 1 arcs	2,107	10,267	11,719
No firsts ed 2 nodes	1,303	6,615	7,567
No firsts ed 2 arcs	4,162	20,418	23,314
No redundancy and 1st ed 2 nodes	1,303	6,615	7,567
No redundancy and 1st ed 2 arcs	4,162	20,418	23,314
Lower order first ed 1 to 2 arcs	6,217	30,569	34,909
Lower order first ed 1 to 2 nodes	1,955	9,923	11,351

Table 3: The sizes of error models as automata

is to establish that simpler error models lead to degraded recall—and not to more generally evaluate the present system as a spell-checker.

The evaluations in this section are performed on quad-core Intel Xeon E5450 running at 3 GHz with 64 GiB of RAM memory. The times are averaged over five test runs of 10,000 words in a stable server environment with no server processes or running graphical interfaces or other uses. The test results are measured using the `getrusage` C function on a system that supports the maximum resident stack size `ru_maxrss` and user time `ru_utime` fields. The times are also verified with the GNU `time` command. The results for `hunspell`, `Voikkospell` and `foma` processes are only measured with `time` and `top`. The respective versions of the software are `Voikkospell 3.3`, `hunspell 1.2.14`, and `Foma 0.9.16alpha`. The reference systems are tested with default settings, meaning that they will only give some fixed number of suggestions whereas our system will calculate all strings within the given error model.

As a reference implementation for English we use `hunspell’s en-US dictionary`⁹ and for a finite-state implementation we use a weighted word-list from [Norvig \(2010\)](http://norvig.com). As a Finnish reference implementation we use `Voikko`¹⁰, with a LAG-based dictionary using `Malaga`¹¹. The reference correction task for Greenlandic is done with `foma’s` ([Huldén, 2009b](http://huldén.com))

⁹<http://wiki.services.openoffice.org/wiki/Dictionaryes>

¹⁰<http://voikko.sf.net>

¹¹<http://home.arcor.de/bjoern-beutel/malaga/>

apply `med` function with default settings¹².

The baseline feature set and the efficiency of spell-checking we are targeting is defined by the currently de facto standard spelling suite in open source systems, hunspell.

In Table 4 we measure the speed of the spell-checking process on native language Wikipedia text with real-world spelling errors and unknown strings. The error model rows are defined as follows: on the *Reference impl.* row, we test the spell-checking speed of the hunspell tool for English, and Voikkospell tool for Finnish. On the *edit distance 2* row we use the basic traditional edit distance 2 without any modifications. On the *No first edits* row we use the error model that does not modify the first character of the word. On the *No redundancy* row we use the edit distance 2 error model with the redundant edit combinations removed. On the *No redundancy and firsts* rows we use the combined error model of *No first edits* and *No redundancy* functionalities. On the row *Lower order first* we apply a lower order edit distance model first, then if no results are found, a higher order model is used. In the tables and formulae we routinely use the language codes to denote the languages: *en* for English, *fi* for Finnish and *kl* for Greenlandic (Kalaallisut).

Error model	En	Fi	Kl
Reference impl.	9.93	7.96	11.42
Generate all edits 2	3818.20	118775.60	36432.80
Edit distance 1	0.26	6.78	4.79
Edit distance 2	7.55	220.42	568.36
No first edits 1	0.44	3.19	3.52
No firsts ed 2	1.38	61.88	386.06
No redundancy ed 2	7.52	4230.94	6420.66
No redundancy and firsts ed 2	1.51	62.05	386.63
Lower order first ed 1 to 2	4.31	157.07	545.91

Table 4: Effect of language and error models to speed (time in seconds per 10,000 word forms)

The results show that not editing the first position does indeed give significant boost to the speed, regardless of language model, which is of course caused by the significant reduction in search space. However, the redundancy avoidance does not seem to make a significant difference. This is most likely because the amount of duplicate paths in the search space is not so proportionally large and their traversal will be relatively fast. The separate application

¹²<http://code.google.com/p/Foma/>

of error models gives the expected timing result between its relevant primary and secondary error models. It should be noteworthy that, when thinking of real world applications, the speed of the most of the models described here is greater than 1 word per second (i.e. 10,000 seconds per 10,000 words).

We measured memory consumption when performing the same tests. Varying the error model had little to no effect. Memory consumption was almost entirely determined by the language model, giving consumptions of 13-7 MiB for English, 0.2 GiB for Finnish and 1.6 GiB for Greenlandic.

To measure the degradation of quality when using different error models we count the proportion of suggestion sets that contain the correct correction among the corrected strings. The suggestion sets are the entire (unrestricted by number) results of correction, with no attempt to evaluate precision¹³. For this test we use automatically generated corpus of spelling errors to get the large-scale results.

Error model	En	Fi	Kl
Edit distance 1	0.89	0.83	0.81
Edit distance 2	0.99	0.95	0.92
Edit distance 3	1.00	0.96	—
No firsts ed 1	0.74	0.73	0.60
No firsts ed 2	0.81	0.82	0.69
No firsts ed 3	0.82	—	—

Table 5: Effect of language and error models to quality (recall, proportion of suggestion sets containing a correctly suggested word)

This test with automatically introduced errors shows us that with uniformly distributed errors the penalty of using an error model that ignores word-initial corrections could be significant. This contrasts to our findings with real world errors, that the distribution of errors tends towards the end of the word, described in 2.2 and (Bhagat, 2007), but it should be noted that degradation can be as bad as given here.

Finally we measure how the text type used will affect the speed of spell-checking. As the best-case scenario we use the unmodified texts of Wikipedia, which contain probably the most realistic native-language-speaker-like typing error dis-

¹³Which, in the absence of suitable error corpora and a more full-fledged language model taking context into account, would be irrelevant for the goal at hand.

tribution available. For text with more errors, where the majority of errors should be recoverable, we introduce automatically generated errors in the Wikipedia texts. Finally to see the performance in the worst case scenario where most of the words have unrecoverable spelling errors we use texts from other languages, in this case English texts for Finnish and Greenlandic spell-checking and Finnish texts for English spell-checking, which should bring us close to the lower bounds on performance. The effects of text type (i.e. frequency of non-words) on speed of spell-checking is given in Table 6. All of the tests in this category were performed with error models under the *avoid redundancy and firsts ed 2* row in previous tables, which gave us the best speed/quality ratio.

Error model	En	Fi	Kl
Native Lang. Corpus	1.38	61.88	386.06
Added automatic errors	6.91	95.01	551.81
Text in another language	22.40	148.86	783.64

Table 6: Effect of text type on error models to speed (in seconds per 10,000 word-forms)

Here we chiefly note that the amount of non-words in text directly reflects the speed of spell-checking. This shows that the dominating factor of the speed of spell-checking is indeed in the correcting of misspelled words.

5 Conclusions and Future Work

In this article, we built a full-fledged finite-state spell-checking system from existing finite-state language models and generated error models. This work uses the system initially described in Pirinen and Lindén (2010) and an algorithm described in Lindén et al. (2012), providing an extensive quantitative evaluation of various combinations of constituents for such a system, and applying it to the most challenging linguistic environments available for testing. We showed that using on-line composition of the word form, error model and dictionary is usable for morphologically complex languages. Furthermore we showed that the error models can be automatically optimised in several ways to gain some speed at cost of recall.

We showed that the memory consumption of the spell-checking process is mainly unaffected by the

selection of error model, apart from the need to store a greater set of suggestions for models that generate more suggestions. The error models may therefore be quite freely changed in real world applications as needed.

We verified that correcting only the first input letter affords a significant speed improvement, but that this improvement is not greatly dependent on the position of such a restriction. This practice is somewhat supported by our tentative finding that it may cause the least drop in practical recall figures, at least in Finnish. It is promising especially in conjunction with a fallback model that does correct the first letter.

We described a way to avoid having a finite-state error model perform redundant work, such as deleting and inserting the same letter in succession. The practical improvement from doing this is extremely modest, and it increases the size of the error model.

In this research we focused on differences in automatically generated error models and their optimisations in the case of morphologically complex languages. For future research we intend to study more realistic error models induced from actual error corpora (e.g. Brill and Moore (2000)). Research into different ways to induce weights into the language models, as well as further use of context in finite-state spell-checking (as in Pirinen et al. (2012)), is warranted.

Acknowledgements

We thank the anonymous reviewers for their comments and the HFST research team for fruitful discussions on the article’s topics. The first author thanks the people of *Oqaasertassualeriffik* for introducing the problems and possibilities of finite-state applications to the morphologically complex language of Greenlandic.

References

- Cyril Allauzen and Mehryar Mohri. 2009. N-way composition of weighted finite-state transducers. *International Journal of Foundations of Computer Science*, 20:613–627.
- Kenneth R Beesley and Lauri Karttunen. 2003. *Finite State Morphology*. CSLI publications.

- Meenu Bhagat. 2007. Spelling error pattern analysis of punjabi typed text. Master's thesis, Thapar University.
- Eric Brill and Robert C. Moore. 2000. An improved error model for noisy channel spelling correction. In *ACL '00: Proceedings of the 38th Annual Meeting on Association for Computational Linguistics*, pages 286–293, Morristown, NJ, USA. Association for Computational Linguistics.
- Fred J Damerau. 1964. A technique for computer detection and correction of spelling errors. *Commun. ACM*, (7).
- Ahmed Hassan, Sara Noeman, and Hany Hassan. 2008. Language independent text correction using finite state automata. In *Proceedings of the Third International Joint Conference on Natural Language Processing*, volume 2, pages 913–918.
- Måns Huldén. 2009a. Fast approximate string matching with finite automata. *Procesamiento del Lenguaje Natural*, 43:57–64.
- Måns Huldén. 2009b. Foma: a finite-state compiler and library. In *Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics: Demonstrations Session*, EACL '09, pages 29–32, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Kimmo Koskenniemi. 1983. *Two-level Morphology: A General Computational Model for Word-Form Recognition and Production*. Ph.D. thesis, University of Helsinki.
- Karen Kukich. 1992. Techniques for automatically correcting words in text. *ACM Comput. Surv.*, 24(4):377–439.
- Vladimir I. Levenshtein. 1966. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics—Doklady 10, 707710. Translated from Doklady Akademii Nauk SSSR*, pages 845–848.
- Krister Lindén, Erik Axelson, Senka Drobac, Sam Hardwick, Miikka Silfverberg, and Tommi A Pirinen. 2012. Using hfst for creating computational linguistic applications. In *Proceedings of Computational Linguistics - Applications, 2012*, page to appear.
- Petar Nikolaev Mitankin. 2005. Universal levenshtein automata. building and properties. Master's thesis, University of Sofia.
- Roger Mitton. 2009. Ordering the suggestions of a spellchecker without using context*. *Nat. Lang. Eng.*, 15(2):173–192.
- Peter Norvig. 2010. How to write a spelling corrector. referred 2011-01-11, available <http://norvig.com/spell-correct.html>.
- Kemal Oflazer. 1996. Error-tolerant finite-state recognition with applications to morphological analysis and spelling correction. *Comput. Linguist.*, 22(1):73–89.
- Tommi A Pirinen and Krister Lindén. 2010. Finite-state spell-checking with weighted language and error models. In *Proceedings of the Seventh SaLTMiL workshop on creation and use of basic lexical resources for less-resourced languages*, pages 13–18, Valletta, Malta.
- Tommi Pirinen, Miikka Silfverberg, and Krister Linden. 2012. Improving finite-state spell-checker suggestions with part of speech n-grams. In *International Journal of Computational Linguistics and Applications IJ-CLA (to appear)*.
- Tommi A Pirinen. 2011. Modularisation of finnish finite-state language description towards wide collaboration in open source development of morphological analyser. In *Proceedings of Nodalida*, volume 18 of *NEALT proceedings*.
- Harri Pitkänen. 2006. Hunspell-in kesäkoodi 2006: Final report. Technical report.
- Joseph J. Pollock and Antonio Zamora. 1984. Automatic spelling correction in scientific and scholarly text. *Commun. ACM*, 27(4):358–368, April.
- Klaus Schulz and Stoyan Mihov. 2002. Fast string correction with levenshtein-automata. *International Journal of Document Analysis and Recognition*, 5:67–85.
- Emmanuel J Yannakoudakis and D Fawthrop. 1983. An intelligent spelling error corrector. *Information Processing and Management*, 19(2):101–108.