

# Predicting Authorship and Author Traits from Keystroke Dynamics

**Barbara Plank**

IT University of Copenhagen  
Department of Computer Science  
Rued Langgaards Vej 7, 2300 Copenhagen S, Denmark  
bapl@itu.dk

## Abstract

Written text transmits a good deal of non-verbal information related to the author’s identity and social factors, such as age, gender and personality. However, it is less known to what extent *behavioral biometric traces* transmit such information. We use typist data to study the predictiveness of authorship, and present first experiments on predicting both age and gender from keystroke dynamics. Our results show that the model based on keystroke features leads to significantly higher accuracies for authorship than the text-based system, while being two orders of magnitude smaller. For user attribute prediction, the best approach is to combine the two, suggesting that extra-linguistic factors are disclosed to a larger degree in written text, while author identity is better transmitted in typing behavior.

## 1 Introduction

Language is a social phenomenon (Nguyen et al., 2015). Whenever we speak or write we transmit a good deal of additional non-verbal information that is related to identity and social factors of an author. Early work in authorship analysis has typically been concerned with finding the author of a text, i.e., *authorship attribution* (Mosteller and Wallace, 1964; Stamatatos, 2009). In recent years, there has been a surge of interest towards the social dimension of language. Studies are interested in linking social factors with linguistic features, e.g., (Eisenstein et al., 2011; Bamman et al., 2014), studying data biases (Hovy and Søgaard, 2015) or building actual attribute prediction models from linguistic features (i.e., *author profiling*). Modeling author traits can further help to improve prediction of related attributes (Liu et al., 2016; Benton et al., 2017), help debiasing models (Hovy, 2015; Zhang et al., 2018) or can be used for a wide range of applications like customer sup-

port, healthcare and personalized machine translation (Mirkin et al., 2015; Rabinovich et al., 2017). Factors studied so far include gender, age, personality or income, to name but a few (Mairesse and Walker, 2006; Luyckx and Daelemans, 2008; Rao et al., 2010; Rosenthal and McKeown, 2011; Nguyen et al., 2011; Volkova et al., 2013; Flekova et al., 2016b; Verhoeven et al., 2016; van Dalen et al., 2017; Ljubešić et al., 2017; Emmerly et al., 2017; van der Goot et al., 2018).

A key question in authorship analysis and profiling is what sorts of *evidence* might bear on determining authorship (Nerbonne, 2007) (or traits). What all prior work has in common is that it *almost exclusively focused on the written text itself*. As people read or write texts, they unconsciously produce cognitive by-product, such as gaze patterns or typist behavior. This evokes and motivates our research question: to what extent is behavioral data beyond the text predictive of authorship and author traits? In this paper we focus on *keystroke dynamics*. They concern a user’s typing pattern. Keystroke logs have the distinct advantage over other cognitive modalities like brain scans or gaze, that keystroke logs are more readily available; they do not rely on special equipment beyond a keyboard. While keystrokes are known to be informative for author verification (cf. Section 5), it is less clear to what extent keystrokes are predictive of authorship, and even more so, of author traits.

**Contributions** a) We study the effect of keystrokes to identify authorship in two corpora of varying size. b) We investigate the predictive power of typist data for age and gender prediction. c) We compare behavioral measures to traditional stylometric features.

## 2 Keystroke dynamics

Keystroke logs are recordings of a user’s typing dynamics. When a person types on a keyboard, the latencies between successive keystrokes and their duration reflect the typing behavior of a person. For example, Figure 2 shows the keystroke hold times (average over single letters) of two users from our dataset. In its raw form, keystroke logs contain information on which key was pressed for how long (key, time press, time release). Research on keystroke dynamics typically consider timing measures derived from *time press* and *time release* events between keystrokes, such as key hold times or interkey durations (see Figure 1).




Figure 1: Keystroke logs illustrated:  $p$  are pauses between keystrokes. Figure adapted from (Goodkind and Rosenberg, 2015).

Only very recently this source has been explored as information in natural language processing, for example, to aid shallow syntactic parsing (Plank, 2016) or deception detection (Banerjee et al., 2014) (Section 5). Keystroke logs have been used in computer security for user verification, however, combining keystroke biometrics with traditional stylometry metrics has not been proven successful (Stewart et al., 2011). The authors focused on a single task and dataset only. In contrast, in this paper we examine to what extent keystroke dynamics are informative for authorship attribution *and* author profiling.

## 3 Experiments

Given a dataset with keystroke logs, we run two sets of experiments: a) authorship attribution, i.e., to determine *who* wrote a given piece of text; and b) authorship profiling, i.e., to determine extralinguistic user traits, in particular age and gender.




Figure 2: Keystroke basic feature set: normalized mean key hold time for two example users.

**Datasets** The two keystroke datasets differ in the amount of users and available meta-data. The first, STEWART, stems from students taking a test on spreadsheet modeling (Stewart et al., 2011). This dataset is not distributed with further meta-data, hence it is used for authorship attribution only. The second dataset, VILLANI (Tappert et al., 2009), is larger (144 participants) and contains demographic meta-data. Keystrokes were recorded for two tasks: free text production and a copy task (fixed text snippet). As we are interested in author attribution/profiling, we consider only the former.

**Pre-processing and Features** First, we remove users with fewer than 5 typing sessions, sessions shorter than 5 words, users without demographics and users that only participated in the copy task (for VILLANI). We also removed two spammers (random skribble). This resulted in a dataset with 34 and 121 users with an average of 99 and 125 tokens per session for STEWART and VILLANI, respectively. The final gender/age distribution is not balanced: 53 female/68 male users, and 56 users above/65 user below thirty. For all keystrokes, the type of key was derived: letters, numbers, punctuation etc., ignoring control keys (FN etc).

Second, we derive 218 biometric features following (Stewart et al., 2011; Tappert et al., 2010). These biometric features include duration features (mean and standard deviation) and are grouped into: i) basic keystroke features, i.e., key hold time (key press and release time) features of the 26 letters from the English alphabet (cf. Figure 2 for an illustration); and ii) extended features: key hold times over groups of keys (like digits, punctuation etc) and transition (inter-key duration) features between successive keystrokes, e.g., between letters and non-letters, or individual letters and

AUTHORSHIP ATTRIBUTION

STEWART (34 users)	#FEATURES	F1-SCORE
Baseline (random)	–	0.4
Text	54k	50.2
Keystrokes (basic)	52	81.0
Keystrokes (extended)	218	<b>90.2</b>
Keystrokes (ext.) + Text	55k	70.2
Keystrokes (ext.) + Embeds	413	<b>91.4</b>
VILLANI (121 users)	#FEATURES	F1-SCORE
Text	46k	67.1
Keystrokes (basic)	52	2.5
Keystrokes (ext.)	218	<b>85.9</b>
Keystrokes (ext.) + Text	46k	71.1
Keystrokes (ext.) + Embeds	413	<b>88.1</b>

Table 1: Results for authorship attribution (34 users), comparison of features (text vs keystrokes), and combined models. Best result in boldface.

groups of such. For these feature measurements, outlier removal and feature standardization is applied (Stewart et al., 2011).

Finally, we extract the final text from the keystroke logging data (employing revisions/backspaces were appropriate). As features we employ those used by the top performing system of the latest PAN author profiling competition (Basile et al., 2017), i.e., word n-grams and character n-grams. N-gram size is tuned on one fold on STEWART, resulting in word unigrams and character 2-3 grams. We also use word embedding features using Polyglot embeddings of 64 dimensions (Al-Rfou et al., 2013), representing text snippets as average embeddings (CBOW) over all tokens (Collobert et al., 2011), enriched with max, sum, standard deviation and embeddings coverage rate. These features worked best on dev.

**Setup** We use a Support Vector Machine (SVM) (Pedregosa et al., 2011) with linear kernel and  $\ell_2$  regularization, similar to the state-of-the-art in author profiling (Flekova et al., 2016a; Basile et al., 2017). We consider a single session of a user as a data instance, and run experiments using 5-fold cross-validation. For author profiling we ensure that all instances of an author end up in the same fold, to not confound profiling with authorship. We report results using weighted F1-score. To ease replicability, all code is released at: <https://github.com/bplank/aat>

## 4 Results

The results of training a classifier to predict the identity of an author are given in Table 1 and Fig-

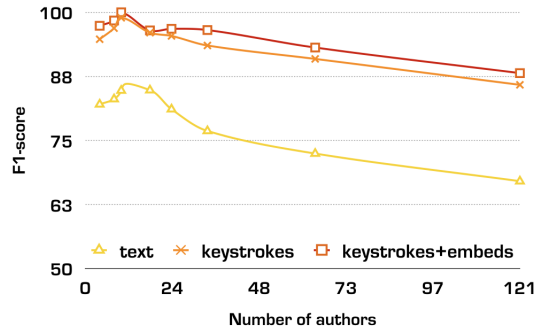


Figure 3: Results for author attribution on VILLANI for increasing number of authors. Keystroke features clearly outperform textual features.

ure 3. The random baseline accuracy is low (0.4% F1). Biometric behavioral features work incredibly well, reaching a performance in the 80-90ies. Already the basic feature set of 52 letter duration features clearly outperforms the stylistic features, reaching 81% F1-score. In contrast, stylometric features from the text alone reach an F1 of only 50%. Note that for the dataset with more users (VILLANI, Figure 3), results for authorship are actually higher, which may be explained by the fact that the smaller dataset is more controlled by topic (exam questions). Figure 3 shows that also on the larger dataset keystroke features outperform the text-based features (word and character n-grams) for authorship, even in setups with few users. These are remarkable results. The behavioral models employ a considerably smaller feature space (cf. column 2 in Table 1). Adding stylometric features improves performance over keystrokes, but only for the embeddings setup, which results in the best setup.

The results for author profiling are given in Table 2. Baseline results (majority baseline) are higher; this task is easier. The gap between stylometric and behavioral features is smaller, but the same trend holds: biometric behavioral features are predictive of gender. To a certain extent this also holds for age (albeit to a lesser extent). Interestingly, combining biometrics with traditional token-based features consistently proves the most effective for author profiling, albeit the best way differs per trait.

Our results suggest that author identity is highly captured in keystrokes alone, while the textual signal provides complementary evidence that together proves the most effective for predicting age and gender of an author.

AUTHOR PROFILING - 121 USERS

	GENDER	AGE	#FEATS
Baseline	33.33	38.06	–
Text	61.34	72.59	80k
Keystrokes (basic)	44.63	37.61	52
Keystrokes (ext.)	63.29	60.58	218
KS + text	60.92	<b>73.25</b>	80k
KS + embeds	<b>63.50</b>	67.12	413

Table 2: Gender and age prediction results, F1-score (age: above/below 30).

## 5 Related Work

Authorship attribution has a long tradition dating back to early works in the 19th century. The most influential work on authorship attribution goes back to [Mosteller and Wallace \(1964\)](#). For a long time approaches to authorship attribution focused on distributions of *function words*, high-frequency words that are presumably not consciously manipulated by the author ([Nerbonne, 2007](#); [Pennebaker, 2011](#)). Recent work also includes authorship studies on microblog texts ([Schwartz et al., 2013](#)). An recent survey is [Stamatatos \(2009\)](#). We here study another source of information that is presumably not consciously manipulated, keystroke dynamics.

A major scientific interest in keystroke dynamics arose in writing research, where it has developed into a promising non-intrusive method for studying cognitive processes involved in writing ([Sullivan et al., 2006](#); [Nottbusch et al., 2007](#); [Wengelin, 2006](#); [Van Waes et al., 2009](#); [Baaijen et al., 2012](#)). In these studies time measurements—pauses, bursts and revisions—are considered traces of the recursive nature of the writing process. *Bursts* are defined as consecutive chunks of text produced and defined by a 2000ms time of inactivity ([Wengelin, 2006](#)). In fact, most prior work that uses keystroke logs focuses on experimental research. For example, [Hanouille et al. \(2015\)](#) study whether a bilingual glossary reduces the working time of professional translators. They consider pause durations before terms extracted from keystroke logs and find that a bilingual glossary reduces the translators’ workload. An analysis of users’ typing behavior was studied by [Baba and Suzuki \(2012\)](#) to measure the impact of spelling mistakes. [Goodkind and Rosenberg \(2015\)](#) investigate pre-word pauses and their re-

lation to multi-word expressions. They found that within MWE pauses vary depending on the cognitive task. [Banerjee et al. \(2014\)](#) were the first to use keystroke patterns for deception detection.

Keystrokes were successfully used for author verification in computer security research ([Stewart et al., 2011](#); [Monaco et al., 2013](#); [Locklear et al., 2014](#)), as they are known to be idiosyncratic ([Leggett and Williams, 1988](#)). Our results show that keystroke biometrics are far superior over stylometry-based features in authorship attribution, and are predictive of author traits.

The study most related to ours ([Stewart et al., 2011](#)) used features from both keystrokes and linguistic stylometry for user verification in a  $k$ -nearest neighbor setup. Their study differs from ours in three aspects. First, they use a more elaborate set of stylometric features (like number of words of a certain length, and readability measures). Second, they target user authentication, thus their setup is a binary classification task (authenticated vs not-authenticated), while we here focus on a multi-class classification setup, which is a considerably more difficult task. Third, they use only a single dataset (STEWART), while we here include results on a second and larger dataset ( $n=121$  authors). To the best of our knowledge, prior work on predicting demographics from typing behavior is typically limited to a single variable ([Tsimperidis et al., 2015](#)), except ([Brizan et al., 2015](#)), whose data is not available. Our study differs from theirs by studying age, and the focus on complementing textual with behavioral data.

**Disclaimer** While modeling user demographics can be seen as one step towards addressing biases in NLP it is important to be aware of potential negative side effects, both from the modeling side through potential exclusion or dual use ([Hovy and Spruit, 2016](#)), as well as the data side, when dealing with privacy sensitive data (cognitive behavioral data) or labels (e.g., mental health).

## 6 Conclusions

We have shown that behavioral biometrics contain highly predictive information for both authorship and author profiling. For authorship attribution, behavioral keystroke metrics significantly outperform traditional text-based features (words and character unigrams), while using a feature set which is orders of magnitude smaller (218 vs sev-

eral thousands of features). In addition, we show that keystroke dynamics are also predictive for author traits (gender and age). Interestingly, for the latter task, it is most beneficial to *combine* behavioral keystroke data with traditional text-based features, suggesting that user *traits* are disclosed to a larger degree in written text while *identity* is better disclosed in typing behavior.

## Acknowledgments

I would like to thank the anonymous reviewers for their valuable feedback. This research is supported by NVIDIA corporation.

## References

- Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2013. *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, chapter Polyglot: Distributed Word Representations for Multilingual NLP. Association for Computational Linguistics.
- Veerle M Baaijen, David Galbraith, and Kees de Glopper. 2012. Keystroke analysis: Reflections on procedures and measures. *Written Communication*.
- Yukino Baba and Hisami Suzuki. 2012. [How are spelling errors generated and corrected? a study of corrected and uncorrected spelling errors using keystroke logs](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 373–377. Association for Computational Linguistics.
- David Bamman, Jacob Eisenstein, and Tyler Schnoebelen. 2014. Gender identity and lexical variation in social media. *Journal of Sociolinguistics*, 18(2):135–160.
- Ritwik Banerjee, Song Feng, Jun Seok Kang, and Yejin Choi. 2014. [Keystroke patterns as prosody in digital writings: A case study with deceptive reviews and essays](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1469–1473, Doha, Qatar. Association for Computational Linguistics.
- Angelo Basile, Gareth Dwyer, Maria Medvedeva, Josine Rawee, Hessel Haagsma, and Malvina Nissim. 2017. N-gram: New groningen author-profiling model. In *Proceedings of the CLEF 2017 Evaluation Labs and Workshop – Working Notes Papers, 11-14 September, Dublin, Ireland (Sept. 2017)*.
- Adrian Benton, Margaret Mitchell, and Dirk Hovy. 2017. [Multi-task learning for mental health using social media text](#). In *Proceedings of European Chapter of Association for Computational Linguistics*. Association for Computational Linguistics.
- David Guy Brizan, Adam Goodkind, Patrick Koch, Kiran Balagani, Vir V Phoha, and Andrew Rosenberg. 2015. Utilizing linguistically enhanced keystroke dynamics to predict typist cognition and demographics. *International Journal of Human-Computer Studies*, 82:57–68.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12(Aug):2493–2537.
- Reinder Gerard van Dalen, Léon Redmar Melein, and Barbara Plank. 2017. Profiling dutch authors on twitter: Discovering political preference and income level. *Computational Linguistics in the Netherlands Journal*, 7:79–92.
- Jacob Eisenstein, Noah A Smith, and Eric P Xing. 2011. Discovering sociolinguistic associations with structured sparsity. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 1365–1374. Association for Computational Linguistics.
- Chris Emmerly, Grzegorz Chrupała, and Walter Daelemans. 2017. Simple queries as distant labels for predicting gender on twitter. In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pages 50–55, Copenhagen, Denmark. Association for Computational Linguistics.
- Lucie Flekova, Jordan Carpenter, Salvatore Giorgi, Lyle Ungar, and Daniel Preotjiuc-Pietro. 2016a. [Analyzing biases in human perception of user age and gender from text](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 843–854, Berlin, Germany. Association for Computational Linguistics.
- Lucie Flekova, Daniel Preotjiuc-Pietro, and Lyle Ungar. 2016b. [Exploring stylistic variation with age and income on twitter](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 313–319. Association for Computational Linguistics.
- Adam Goodkind and Andrew Rosenberg. 2015. [Muddying the multiword expression waters: How cognitive demand affects multiword expression production](#). In *Proceedings of the 11th Workshop on Multiword Expressions*, pages 87–95, Denver, Colorado. Association for Computational Linguistics.
- Rob van der Goot, Nikola Ljubešić, Ian Matroos, Malvina Nissim, and Barbara Plank. 2018. Bleaching text: Abstract features for cross-lingual gender prediction. In *ACL*, Melbourne, Australia. Association for Computational Linguistics.
- Sabien Hanouille, Véronique Hoste, and Aline Remael. 2015. The translation of documentaries: Can do-

- mainspecific, bilingual glossaries reduce the translators workload? an experiment involving professional translators. *New Voices in Translation Studies*, 23:13.
- Dirk Hovy. 2015. [Demographic factors improve classification performance](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 752–762. Association for Computational Linguistics.
- Dirk Hovy and Anders Søgaard. 2015. [Tagging performance correlates with author age](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 483–488, Beijing, China. Association for Computational Linguistics.
- Dirk Hovy and L. Shannon Spruit. 2016. [The social impact of natural language processing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598. Association for Computational Linguistics.
- John Leggett and Glen Williams. 1988. Verifying identity via keystroke characteristics. *International Journal of Man-Machine Studies*, 28(1):67–76.
- Fei Liu, Julien Perez, and Scott Nowson. 2016. A recurrent and compositional model for personality trait recognition from short texts. In *Proceedings of the Workshop on Computational Modeling of People’s Opinions, Personality, and Emotions in Social Media (PEOPLES)*, Osaka, Japan. The COLING 2016 Organizing Committee.
- Nikola Ljubešić, Darja Fišer, and Tomaž Erjavec. 2017. Language-independent gender prediction on twitter. In *Proceedings of the Second Workshop on NLP and Computational Social Science*, pages 1–6, Vancouver, Canada. Association for Computational Linguistics.
- Hilbert Locklear, Sathya Govindarajan, Zdenka Sitova, Adam Goodkind, David Guy Brizan, Andrew Rosenberg, Vir V Phoha, Paolo Gasti, and Kiran S Balagani. 2014. Continuous authentication with cognition-centric text production and revision features. In *Biometrics (IJCB), 2014 IEEE International Joint Conference on*.
- Kim Luyckx and Walter Daelemans. 2008. [Personae: a corpus for author and personality prediction from text](#). In *LREC 2008*.
- François Mairesse and Marilyn Walker. 2006. [Automatic recognition of personality in conversation](#). In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*.
- Shachar Mirkin, Scott Nowson, Caroline Brun, and Julien Perez. 2015. [Motivating personality-aware machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1102–1108, Lisbon, Portugal. Association for Computational Linguistics.
- John V Monaco, John C Stewart, Sung-Hyuk Cha, and Charles C Tappert. 2013. Behavioral biometric verification of student identity in online course assessment and authentication of authors in literary works. In *Biometrics: Theory, Applications and Systems (BTAS), 2013 IEEE Sixth International Conference on*.
- Frederick Mosteller and David Wallace. 1964. *Inference and disputed authorship: The Federalist*. Addison-Wesley.
- John Nerbonne. 2007. The exact analysis of text. *Foreword to the 3rd edition of Frederick Mosteller and David Wallace Inference and Disputed Authorship: The Federalist Papers CSLI: Stanford*.
- Dong Nguyen, A. Seza Dogru, Carolyn Penstein Ros, and Franciska de Jong. 2015. Computational sociolinguistics: A survey. *Computational Linguistics*, 42(3):537–59.
- Dong Nguyen, A. Noah Smith, and P. Carolyn Rosè. 2011. *Proceedings of the 5th ACL-HLT Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, chapter Author Age Prediction from Text using Linear Regression. Association for Computational Linguistics.
- Guido Nottbusch, Rüdiger Weingarten, and Said Sahel. 2007. From written word to written sentence production. *Writing and cognition: Research and applications.*, pages 31–54.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- James W Pennebaker. 2011. Using computer analyses to identify language style and aggressive intent: The secret life of function words. *Dynamics of Asymmetric Conflict*, 4(2):92–102.
- Barbara Plank. 2016. [Keystroke dynamics as signal for shallow syntactic parsing](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 609–619, Osaka, Japan. The COLING 2016 Organizing Committee.
- Barbara Plank and Dirk Hovy. 2015. [Personality traits on twitter—or-how to get 1,500 personality tests in a week](#). In *Proceedings of the 6th Workshop on*

- Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 92–98, Lisboa, Portugal. Association for Computational Linguistics.
- Ella Rabinovich, Raj Nath Patel, Shachar Mirkin, Lucia Specia, and Shuly Wintner. 2017. Personalized machine translation: Preserving original author traits. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 1074–1084, Valencia, Spain. Association for Computational Linguistics.
- Delip Rao, David Yarowsky, Abhishek Shreevats, and Manaswi Gupta. 2010. Classifying latent user attributes in twitter. In *Proceedings of the 2nd international workshop on Search and mining user-generated contents*, pages 37–44. ACM.
- Sara Rosenthal and Kathleen McKeown. 2011. [Age prediction in blogs: A study of style, content, and online behavior in pre- and post-social media generations](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 763–772. Association for Computational Linguistics.
- Roy Schwartz, Oren Tsur, Ari Rappoport, and Moshe Koppel. 2013. [Authorship attribution of micro-messages](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1880–1891. Association for Computational Linguistics.
- Efstathios Stamatatos. 2009. A survey of modern authorship attribution methods. *Journal of the American Society for information Science and Technology*, 60(3):538–556.
- John C Stewart, John V Monaco, Sung-Hyuk Cha, and Charles C Tappert. 2011. An investigation of keystroke and stylometry traits for authenticating online test takers. In *Biometrics (IJCB), 2011 International Joint Conference on*.
- Kirk PH Sullivan, Eva Lindgren, et al. 2006. *Computer keystroke logging and writing: Methods and applications*. Elsevier.
- Charles C Tappert, Sung-Hyuk Cha, Mary Villani, and Robert S Zack. 2010. A keystroke biometric system for long-text input. *International Journal of Information Security and Privacy (IJISP)*, 4:32–60.
- Charles C Tappert, Mary Villani, and Sung-Hyuk Cha. 2009. Keystroke biometric identification and authentication on long-text input. *Behavioral biometrics for human identification: Intelligent applications*, pages 342–367.
- Ioannis Tsimperidis, Vasilios Katos, and Nathan Clarke. 2015. Language-independent gender identification through keystroke analysis. *Information & Computer Security*, 23(3):286–301.
- Luuk Van Waes, Mariëlle Leijten, and Daphne Van Weijen. 2009. Keystroke logging in writing research: Observing writing processes with inputlog. *GFL-German as a foreign language*, 2(3):41–64.
- Ben Verhoeven, Walter Daelemans, and Barbara Plank. 2016. Twisty: A multilingual twitter stylometry corpus for gender and personality profiling. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France. European Language Resources Association (ELRA).
- Svitlana Volkova, Theresa Wilson, and David Yarowsky. 2013. [Exploring demographic language variations to improve multilingual sentiment analysis in social media](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1815–1827. Association for Computational Linguistics.
- Åsa Wengelin. 2006. Examining pauses in writing: Theory, methods and empirical data. *Computer keystroke logging and writing: methods and applications (Studies in Writing)*, 18:107–130.
- Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. 2018. Mitigating unwanted biases with adversarial learning. *arXiv preprint arXiv:1801.07593*.