Department	School of Computing
Supervisors	Peter J Barclay
Project Title	Characterising Automatically Generated Text

PROJECT DESCRIPTION
Research is needed to understand the differences between human and machine generated text. Prior research has focused on identification of identification of deceptive text, such as phishing attempts or bot-generated tweets. There is, however, less work on identifying other forms of generated text such as automatic translations, or text that has been recorded by 'essay assistant' software. This project would focus on characterising the differences between human generated and machine generated text, for example by comparing originally authored material with automatically rewritten versions.

Perspective applicants are encouraged to contact the Supervisor before submitting their applications. Applications should make it clear the project you are applying for and the name of the supervisors.

Academic qualifications
A first degree (at least a 2.1) ideally in Computing or a related discipline, with a good fundamental knowledge of natural language processing.

English language requirement
IELTS score must be at least 6.5 (with not less than 6.0 in each of the four components). Other, equivalent qualifications will be accepted. Full details of the University's policy are available online.

Essential attributes:
- Experience of fundamental machine learning models
- Competent in Python programming
- Knowledge of linguistics or natural language processing
- Good written and oral communication skills
- Strong motivation, with evidence of independent research skills relevant to the project
- Good time management

Desirable attributes:

Indicative Bibliography	Afroz, S., Brennan, M., & Greenstadt, R. (2012). Detecting hoaxes, frauds, and deception in writing style online. *2012 IEEE Symposium on Security and Privacy*, 461–475. Boudin, F., Mougard, H., & Cram, D. (2016). How document pre-processing affects keyphrase extraction performance. *ArXiv Preprint ArXiv:1610.07809*. Dou, Y., Forbes, M., Koncel-Kedziorski, R., Smith, N., & Choi, Y. (2022). Is GPT-3 Text Indistinguishable from Human Text? Scarecrow: A Framework for Scrutinizing Machine Text. *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 7250–7274. https://doi.org/10.18653/v1/2022.acl-long.501

	Hancock, J. T., Curry, L. E., Goorha, S., & Woodworth, M. (2007). On lying and being lied to: A linguistic analysis of deception in computer-mediated communication. *Discourse Processes*, *45*(1), 1–23. Ippolito, D., Duckworth, D., Callison-Burch, C., & Eck, D. (2019). Automatic detection of generated text is easiest when humans are fooled. *ArXiv Preprint ArXiv:1911.00650*. Jawahar, G., Abdul-Mageed, M., & Lakshmanan, L. V. (2020). Automatic detection of machine generated text: A critical survey. *ArXiv Preprint ArXiv:2011.01314*. Newman, M. L., Pennebaker, J. W., Berry, D. S., & Richards, J. M. (2003). Lying words: Predicting deception from linguistic styles. *Personality and Social Psychology Bulletin*, *29*(5), 665–675. Varshney, L. R., Keskar, N. S., & Socher, R. (2020). Limits of detecting text generated by large-scale language models. *2020 Information Theory and Applications Workshop (ITA)*, 1–5.
Enquiries	For informal enquiries about this PhD project, please contact p.barclay@napier.ac.uk
Web page	https://www.napier.ac.uk/research-and-innovation/research-degrees/application-process