

Towards a Generalised Framework for Behaviour Insight Mining*

Allmin Susaiyah^{1,3}[0000-0002-3926-6653], Aki Härmä¹[0000-0002-2966-3305], and
Ehud Reiter²[0000-0002-7548-9504] Rim Helaoui¹[0000-0001-6915-8920] Milan
Petković^{1,3}[0000-0002-6264-6822]

¹ Philips Research, High Tech Campus 34, Eindhoven, Netherlands

² University of Aberdeen, Aberdeen, Scotland

³ Eindhoven University of Technology, Eindhoven, Netherlands

Abstract. A behaviour insight is a statement that explains the behaviour of a user's activity or a system's performance under certain conditions. Care must be taken to deliver insights in a way that is useful to the user. Mining for such insights and generating text to communicate about those is a complicated process and there is a need for standardisation to ensure it is done properly. In this work, we lay down the foundation towards a generalised behaviour insight mining framework based on our previous experiences in developing them and express our views on how it can be adopted in a wide range of applications such as Sleep, Marketing, Medical Image Diagnosis and clinical workflow management.

Keywords: data-to-text system · behaviour-insight-mining · behaviour-change-support-system

1 Introduction

Over the last two years alone, 90 percent of the data in the world was generated. Today, a human on average generates more data than the entire global data store that was available just a few decades ago. Thanks to Social Media, the Internet of Things, Digital Photography and Internet services, we now have more user data than ever. This data contains a large amount of useful information that can help to improve behaviour, such as sleep, diet, education, driving, and even specialised tasks such as marketing campaigns.

There have been several behaviour change support systems (BCSS), that generate text based on the user behaviour learned from data. The key objective of such systems is to understand user behaviour and persuade them to improve it. Such systems can be considered as a natural language generation system or otherwise data-to-text systems catering to the specific use-case of behaviour change support.

Some of the behaviour change support systems are:

* Supported by Horizon H2020 - PhilHumans Network.

1. Safer Drive [2]: It is a mobile application that generates weekly reports for automobile drivers with an intent to improve their driving habits.
2. Quit Smoking [8]: This system generates smoking cessation letters to encourage users to quit smoking. However, the authors also stress a fact that not always a behavior change support system can have a positive impact on the user.
3. Anselma et al [1] develop a diet BCSS that aims towards persuading the user to follow a healthy diet.

Reiter [7] introduced the data-to-text pipeline that consists of Signal Analysis, Data Interpretation, Document Planning, Microplanning and Realisation. This is widely adopted in many data-to-text systems such as Babytalk [7], SUMTIME [9], SUMTIME-TURBINE [11], SCUBATEXT[10], and even generalise to data-to-text systems that existed long back such as ANA [5].

Behaviour Insight Mining (BIM) systems are data-to-text systems similar to the BCSS, but with the main objective of making the users understand their behavior rather than persuading them directly. With the advent of smartphones, wearable devices, and the internet of things, the interest in BIM systems has been growing rapidly. In Healthcare, BIM systems have been growing rapidly. Hommes et al[4], introduce a personalised quality-of-life support tool for cancer patients. Pauws [6]

The BIM systems require specialised mechanisms for interacting with the user in a relate-able, flexible and fast manner. Hence, we modify the existing data-to-text pipeline and additionally include what we call "insight intelligence" that makes the system flexible to the continuously changing user's needs. In this paper, we show, that the proposed pipeline can also be extended to solve many other problems of Data Mining apart from BIM.

2 Behaviour Insight Mining Pipeline

In this section, we explain the different components of our proposed behaviour insight pipeline as shown in Fig1. Although the figure appears like a feedback loop, the feedback does not happen continuously. It rather occurs irregularly. When feedback happens, the insight intelligence mechanism updates. Hence the system can be considered as a pipeline. The pipeline is greatly inspired from Reiter's data-to-text Pipeline [7]. Similar to Reiter's pipeline we have certain blocks that are dedicated to looking at data and few blocks dedicated to generating text. The key difference is that we introduce an insight intelligence mechanism that gives our approach the flexibility to incorporate user preferences as will be discussed in the upcoming sections.

The data should contain all the fields that are required to generate all types of insights that the user wants currently or might want in the future. The data is also governed by the technical feasibility of the system. The data is collected using many sensors, it can be one or many of the following: measuring device, camera, activity loggers, etc. This data can be of following, not limited to, generic types:

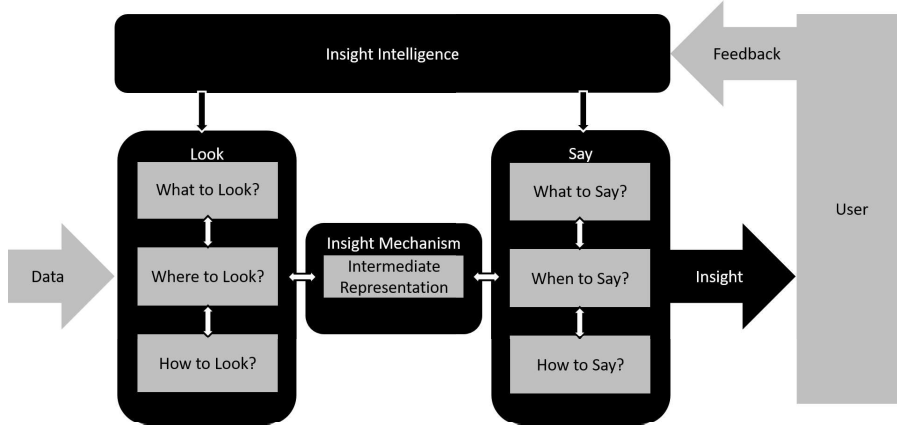


Fig. 1. The Proposed pipeline for behaviour insight mining.

- Metadata: It contains some information about the data itself. Mostly useful in identifying and referencing the data.
- Timestamps: This helps to associate events to the time of occurrence.
- Measurements: This is the most important data that is analysed. It is a quantity of some measure of the user’s behaviour expressed in terms of a consistent unit.

2.1 What-to-look

The what-to-look block defines the data that must be analysed to generate insights. It mainly depends upon the domain that the pipeline is deployed, the target user and type of data. If the pipeline is used for brushing behaviour, then the system should look at brushing duration, percentage of brushing, etc. If used for medical report generation, the system should look at scans and supporting information. More importantly, in Europe, General Data Protection Regulation regulations must be taken into consideration at this stage.

2.2 Where-to-look

Generally, the size of data stored per user is huge and the time to generate insights is limited. This block defines wherein the data, the insight mechanism should look to discover insights. This prevents the generator from looking into irrelevant part of the data. For example, for radiology applications, it would be the insight mechanism should look inside the relevant anatomy in the scans. Sometimes, the data might be missing due to many reasons such as malfunctioning sensor, transmission failure etc. During this scenario, the where-to-look module must be aware of what needs to be done as a workaround at this situation. The insight intelligence component could contain such workarounds.

2.3 How-to-look

The data holds important relationships that need to be considered while looking at it. The way we look at time data will differ from the way we look at rank data. Different data operators would be applicable to different types of data. For example, computing the average time to bed, given the time to bed as timestamps is non-trivial as compared to computing the average sleep duration. At the end of this pipeline, we have insights about the user behaviour that will be contrasting and significantly different [3] to other insights.

2.4 Intermediate Representation

The important insights are stored as an intermediate representation. It is similar to the "messages" as explained in [7]. The intermediate representation could be structured as semantic triples, This makes it simpler for the saying challenges (explained below). The triples have a subject, an object and a predicate that relates the subject to the object. There could also be complex triples where either or both subjects and objects can be triples themselves. A sample triple is as follows: subject(John), predicate(has less sleep on), object(Wednesdays).

2.5 What-to-say

The what-to-say challenge is similar to the "content determination" stage in Reiter's pipeline [7]. Certain information is more important to the user and this changes with time. Hence, there should be a mechanism to determine the insights to convey to the user. Apart from identifying important insights, this system should as try to avoid dangerous insights. For example, "you are likely to get a heart attack in the next 5 minutes" is an insight that should be avoided.

2.6 When-to-say

Usually, the users require multiple insights and the order in which it should be conveyed is important. It depends on the application and the user's preferences. There is also a need for consideration when the insight needs to be delivered. For example, the insights should not be delivered to the user when he is in a busy meeting. The when-to-say component sorts the the insights in the most appropriate order and delivers it to the user at the appropriate time. It is similar to the text structuring explained in [7].

2.7 How-to-say

It is similar to the "Microplanning" module in [7]. The way the insight needs to be phrased is very important and needs to consider the following:

1. Psychology: The how-to-say must be careful in phrasing the words in order to not sound patronising or demotivating.

2. Confidence: The uncertainties in the data must be considered and mentioned while communicating the insights with the user.
3. Language Difficulty: Not all users understand the same level of language. this should be taken into consideration while realising the text.

2.8 User Feedback

Personalising the insight experience is the most important aspect of a Behaviour insight mining system. Hence, user feedback must be obtained and used to deliver relevant and interesting insights.

2.9 Insight Intelligence

Insight Intelligence stores all the information about the user's preferences that would help the looking and saying blocks. This is the governing block that sends information to other blocks such that the blocks are aware of user preferences. For example, it consists of a record of how the user's interests have been varying over time and help the What-to-say block. It also has specific domain knowledge that helps the how-to-look component. If the user starts to log additional data or is interested in a new type of insight, the insight intelligence communicates to the what-to-look block and how-to-look blocks appropriately. This block is also aware of the user's context such as, being in the office, jogging, driving, etc. Hence this helps the when-to-say block to consider this information. The insight intelligence is also capable to discern user's understanding of the insight and thereby can trigger modifications to the language generation model of the how-to-say block.

3 Pipeline Adaptation for Different Domains

In this section, we present how the proposed pipeline can be applied for different application such as Sleep Behaviour Change Support Systems, Marketing Support, Medical Image Diagnosis Report Generation and Clinical workflow management

3.1 Sleep Behaviour Change Support Systems

The quality of sleep is important for healthy living and improving it is imperative. By understanding, sleeping behaviour, we bring a positive change towards sleep behaviour. The what-to-look block reads the user data and determines parameters such as sleep duration, sleep latency and sleep interruption for a given user in a given period of time. The where-to-look looks for the above parameters across different contexts such as during weekdays, weekends, workdays, holidays, etc. The how-to-look part treats different measures differently while performing a comparison. For example, for a night of healthy sleep, the sleep duration should ideally be more (up to 8 hours), whereas, sleep latency must be

as less as possible. The how-to-look component takes care of this with the help of domain knowledge of the measurement. During the initial days of deployment, in the absence of a user model, the BIM system presents statistically significant insights to the user. The significance is measured between the same measurement in two different contexts. The intermediate representation stores all significant insights. The what-to-say block determines the insights that are interesting to the user with the help of an insight intelligence block. The sleep insights are best to be told after bedtime, to give immediate feedback to explain the quality of sleep to the user. This is taken care of by the when-to-say block. The how-to-say block contains templates at various levels of language complexity to express the qualifying insights to the user. The user gives his feedback on:

1. Relevance of the Insight: This helps the insight intelligence to promote more relevant insights
2. Quality of wording: this allows the insight intelligence to adjust the language complexity

3.2 Marketing Support

This is an example of how a system's behaviour can be analysed using our pipeline. Marketing through campaigns requires making decisions as to which campaign is received well in which region. These decisions help the marketing manager to improve the marketing strategy, for example, by promoting campaigns that are being received well and cutting down campaigns that are not. Generally, there can be millions of insights arising from the data. Standardising and automating this process saves the company millions of dollars each year. BIM can help the Marketing managers by generating important insights (out of millions of possible insights) from the marketing data and support them in their decision-making process. The main difference between this use case from the previous one is that the source of data and the end-user are not the same. Still, the proposed pipeline can be implemented. The what-to-look part looks at the key performance indicators of various campaigns and products in various countries. The where-to-look imposes certain domain constraints on the data. This helps the BIM system to avoid unnecessary insights. For example, the marketing campaign in China does not have to be compared with the marketing campaigns in Brazil. The how-to-look part defines how each KPI needs to be evaluated as some of them need to be high and some low. The what-to-say part focuses on a mix of insights indicating campaigns that are improving and failing. There could also be a different mix of preferences that the insight intelligence system could learn from the marketing manager's feedback and incorporate. The when-to-say is crucial too as different regions have different time zones. The insights should be generated during the low activity period(night) for each region such that it encompasses each day's activity as a whole. Finally, how-to-say part decides how the sentences depicting campaign performances need to be framed. It should incorporate sentence variation, confidence, trace-ability and transparency. The

insight intelligence helps to adapt the system to the likes and needs of the marketing manager by knowing what type of insights are interesting, which region is focused more, etc.

3.3 Medical Image Diagnosis Report Generation and Clinical workflow management

This is an example of how a combination of system and user behaviour is analysed for behaviour insight mining. A simple use case is the medical diagnosis of lung cancer. The what-to-look part examines chest Computed Tomography or Magnetic Resonance Imaging, The where-to-look part segments the scans for lesions and the how-to-look part looks at each lesion and computes features that indicate abnormalities. Here, the contexts for comparison can be different lesions captured from the same patient or same lesion of the same patient captured at different time periods. The what-to-say part determines the insights about lesions that are interesting to the radiologist. The when-to-say block determines which patients and lesions are crucial and helps in clinical triage. The how-to-say part performs document planning as described by Reiter [7] to present the findings to the radiologist in a suitable format. The insight intelligence mechanism learns the radiologist's behaviour and preferences and adjusts each individual block.

4 Conclusion

In this work, we propose the movement of behaviour insight mining (BIM) systems towards a generalised pipeline that slightly differs from the existing data-to-text pipeline. We believe that the proposed approach could solve many of the challenges in BIM systems and explain how it is implemented for several BIM applications such as sleep behaviour analysis, marketing support, Medical Image Diagnosis, and clinical workflow management. Since behaviour insight mining is a field that is at its peak, due to the surge of personal fitness trackers and connected devices, the proposed pipeline is relevant to the current scenario.

5 Acknowledgement

This work was supported by the Horizon H2020 Marie Skłodowska-Curie Actions Initial Training Network European Industrial Doctorates project under grant agreement No. 812882 (PhilHumans). The authors would like to acknowledge the CLAN reading group at the University of Aberdeen for their valuable feedback and discussions that greatly helped this work.

References

1. Anselma, L., Mazzei, A.: Designing and testing the messages produced by a virtual dietitian. In: Proceedings of the 11th International Conference on Natural Language Generation. pp. 244–253. Association for Computational Linguistics, Tilburg

- University, The Netherlands (Nov 2018). <https://doi.org/10.18653/v1/W18-6531>, <https://www.aclweb.org/anthology/W18-6531>
2. Braun, D., Reiter, E., Siddharthan, A.: Saferdrive: An nlg-based behaviour change support system for drivers. *Natural Language Engineering* **24**(4), 551–588 (2018)
 3. Härmä, A., Helaoui, R.: Probabilistic scoring of validated insights for personal health services. In: 2016 IEEE Symposium Series on Computational Intelligence (SSCI). pp. 1–6. IEEE (2016)
 4. Hommes, S., van der Lee, C., Clouth, F., Vermunt, J., Verbeek, X., Krahrmer, E.: A personalized data-to-text support tool for cancer patients. In: Proceedings of the 12th International Conference on Natural Language Generation. pp. 443–452. Association for Computational Linguistics, Tokyo, Japan (Oct–Nov 2019). <https://doi.org/10.18653/v1/W19-8656>, <https://www.aclweb.org/anthology/W19-8656>
 5. Kukich, K.: Design of a knowledge-based report generator. In: Proceedings of the 21st annual meeting on Association for Computational Linguistics. pp. 145–150. Association for Computational Linguistics (1983)
 6. Pauws, S., Gatt, A., Krahrmer, E., Reiter, E.: Making effective use of healthcare data using data-to-text technology. *CoRR* **abs/1808.03507** (2018), <http://arxiv.org/abs/1808.03507>
 7. Reiter, E.: An architecture for data-to-text systems. In: Proceedings of the Eleventh European Workshop on Natural Language Generation. pp. 97–104. ENLG '07, Association for Computational Linguistics, Stroudsburg, PA, USA (2007), <http://dl.acm.org/citation.cfm?id=1610163.1610180>
 8. Reiter, E., Robertson, R., Osman, L.M.: Lessons from a failure: Generating tailored smoking cessation letters. *Artificial Intelligence* **144**(1-2), 41–58 (2003)
 9. Sripada, S., Reiter, E., Davy, I.: Sumtime-mousam: Configurable marine weather forecast generator. *Expert Update* **6**(3), 4–10 (2003)
 10. Sripada, S.G., Gao, F.: Linguistic interpretations of scuba dive computer data. In: 2007 11th International Conference Information Visualization (IV'07). pp. 436–441. IEEE (2007)
 11. Yu, J., Reiter, E., Hunter, J., Sripada, S.: Sumtime-turbine: a knowledge-based system to communicate gas turbine time-series data. In: International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems. pp. 379–384. Springer (2003)