

# Schema-Driven Actionable Insight Generation and Smart Recommendation

Citation for published version (APA):

Susaiyah, A., Härmä, A., & Petković, M. (2023). *Schema-Driven Actionable Insight Generation and Smart Recommendation*.

## Document status and date:

Published: 24/07/2023

## Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

## General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

[www.umlib.nl/taverne-license](http://www.umlib.nl/taverne-license)

## Take down policy

If you believe that this document breaches copyright please contact us at:

[repository@maastrichtuniversity.nl](mailto:repository@maastrichtuniversity.nl)

providing details and we will investigate your claim.

---

# SCHEMA-DRIVEN ACTIONABLE INSIGHT GENERATION AND SMART RECOMMENDATION

---

A PREPRINT

✉ **Allmin Susaiyah**

Dept. of Mathematics and Computer Science  
Eindhoven University of Technology  
Netherlands  
a.p.s.susaiyah@tue.nl

✉ **Aki Härmä**

Advanced Computing Sciences  
Maastricht University  
The Netherlands  
aki.harma@maastrichtuniversity.nl

✉ **Milan Petković**

Dept. of Mathematics and Computer Science  
Eindhoven University of Technology  
The Netherlands  
m.petkovic@tue.nl

July 26, 2023

## ABSTRACT

In natural language generation (NLG), insight mining is seen as a data-to-text task, where data is mined for interesting patterns and verbalised into 'insight' statements. An 'over-generate and rank' paradigm is intuitively used to generate such insights. The multidimensionality and subjectivity of this process make it challenging. This paper introduces a schema-driven method to generate actionable insights from data to drive growth and change. It also introduces a technique to rank the insights to align with user interests based on their feedback. We show preliminary qualitative results of the insights generated using our technique and demonstrate its ability to adapt to feedback.

**Keywords** Schema · Insight Generation · Insight Recommendation · Neural Networks

## 1 Introduction

An insight is a deep and accurate comprehension of behaviour or patterns observed in data. An actionable insight helps to drive growth and change Jørnø and Gynther [2018]. They can be expressed as a comparison of a measurement made in two distinct yet equivalent contexts, as in Härmä and Helaoui [2016]. For example, one could compare a user's sleep measurements between two comparable contexts say during Mondays and other days. This can be written as, "On Mondays, you sleep less than on the other days". Such insights when communicated to the user help them to understand Reiter [2007] or take corrective actions to their behaviour Jørnø and Gynther [2018], O'Malley et al. [2012].

There are three important criteria for an actionable insight: 1) truthfulness Funke et al. [2018], Agrawal and Shafer [1996], Härmä and Helaoui [2016] 2) significance Agrawal and Shafer [1996], Härmä and Helaoui [2016] and 3) usefulness Freitas [1999], Fayyad et al. [1996]. The truthfulness of an insight can be asserted using statistical tests and the significance can be asserted using mathematical models. Meanwhile, the usefulness of insight depends on various factors like quantify-ability, meaningfulness, understandability and surprisingness Freitas [1999], Fayyad et al. [1996], Susaiyah et al. [2020a], Sudarsanam et al. [2019].

The state-of-the-art natural language generation models such as GPTs Radford et al. [2019], Brown et al. [2020] have limitations when generating insights from large datasets due to their size constraints and hallucinations. Generating actionable insights on a large scale can be done using the generalised framework (GENF) introduced in Susaiyah et al. [2020b] and the 'over-generate and rank' (OGR) paradigm Gatt and Kraemer [2018]. The GENF framework talks

Schema Definition				
template		applicable_items	scoring type	
<code>{{age_split:1}} ({{mean:1}}) the ({{measurement}})</code> <code>{{tense(be)}} ({{comparison}}) ({{!age_split:2}}) ({{mean:2}})</code>		<code>{'measurement':['requested_dose',</code> <code>'acquisition_time','exam_duration']}</code>	ks	
Measurement Definition				
measurement	column_in_data	phrase	regular comparative phrase	
requested_dose	requested_dose	requested dose	['lower than', 'equal to', 'greater than']	
Context Definition				
name	entities	query	prepositional phrase	pair id
young	['younger than 30']	["data.age_upper_limit<=30"]	['in patients younger than 30 years']	1
middle	['30-50']	["(data.age_lower_limit>=30) & (data.age_upper_limit<=50)"]	['in patients of age 30 to 50']	1

**Actionable Insight**

In patients younger than 30 years (164.77 units) the requested dose is greater than in patients of other age groups (45.30 units)

Figure 1: Insight schema definition and a sample insight. The colors indicate corresponding blocks

about components that look at (analyse) the data, represent the insights in an intermediary format, say (generate) the statement, obtain user feedback and upgrade the system appropriately. Once many truthful and significant insights are generated using OGR, recommendation algorithms can be used to rank them and select useful insights.

This paper introduces the schema-driven actionable insight-generation approach that grew out of our recent experiences in various applications. The main contribution of this paper is the method to generate actionable insights on a large scale using a controlled natural language input that we call an insight schema. It is a direct implementation of the GENF framework. As a secondary contribution, we present a training protocol that can be used by machine learning models to recommend relevant and diverse insights to the user even with very little feedback. Thirdly, GEN-IG<sup>1</sup>: A toolkit to perform this along with full documentation is available in <https://github.com/allmin/gen-ig>.

## 2 Insight schema

The insight schemas allow us to define prospective insights using a controlled language and generate them complying with the actionability constraints. We developed the GEN-IG toolkit to understand such schemas and generate insights.

### 2.1 Insight schema definition

The 'template' in the schema definition shown in Figure 1 mentions the data necessary to generate the insights along with the technique to validate them in 'scoring type'. The template compares a continuous variable: 'measurement' across two age groups (categorical variables): 'age\_split:1' and 'age\_split:2'. The eligible 'measurement' for the schema is enumerated in the column 'applicable\_items'. The schema includes mean:1 and mean:2 to display the mean of the 'measurement' in contexts 1 and 2 respectively. This helps in the actionability of such insights. An insight belonging to the above schema is also shown in the figure. The recommended "scoring tests" applicable to different types of schema templates are given below:

1. Two-sample Kolmogorov-Smirnov (KS) test Naaman [2021] and the Mann-Whitney U test Fay and Proschan [2010]: Used while comparing the distributions of measurement across two contexts.
2. Binomial test: Used when one of the contexts is a scalar or when the comparison is made on the frequencies of the contexts instead of their distributions.

Other types of statistical tests can be used with our approach depending on what characteristics of the data are being compared.

### 2.2 Measurement and context definition

In addition to the insight schema definitions, the measurement and context definitions help with text realisation and data querying respectively. Figure 1 shows the measurement definition. The highlighted definition realises "measurement tense(be) comparison" of Figure 1 into: "the requested dose is greater than" where the verb 'be' takes the right tense and subject-verb agreement based on the contexts and subject-verb agreement. Additionally, a tolerance level  $\tau$  is defined which is used to rank the insights (see Section 3.2).

<sup>1</sup><https://youtu.be/dAzHRPT1oUg>

Figure 1 also shows a sample context definition that specifies, for each contexts, the pandas pandas development team [2020] or SQL queries to extract the data. When two or more contexts are comparable, for example, 'in patients younger than 30 years' and 'in patients of age 30 to 50', they are assigned the same pair id.

### 3 Stages of insight generation

The stages of the proposed approach are elaborated below.

#### 3.1 Insight library generation

It corresponds to the what, where, and how-to-look components of the GENF. Here, individual measurements and contexts are enumerated based on the insight schema definitions. The enumerations are used to create intermediate templates and queries that will be used in the subsequent steps.

#### 3.2 Insight scoring

Scoring the insights helps in recommending them. Each insight is tested for its truthfulness using statistical tests mentioned in Section 2.1. Subsequently, the true insights ( $p < 0.05$ ) are assigned a relevance score based on their completeness, significance defined by the tolerance level  $\tau$  and usefulness. The completeness score  $Score_C$  as shown in Equation 1 is calculated using the sampling rate of measurements  $F_{exp}$ , the time-span of the queried data  $T$ , and the actual recorded data samples  $N_{rec}$ . The significance score  $Score_S$  is calculated using  $\gamma$ , which determines the margin of error and  $\delta$ , the calculated difference of means of the measurements across the two contexts.

A neural network is used to provide a usefulness score  $Score_U$  to the insights. For this, user feedbacks: "not useful at all", "not useful", "neutral", "useful", "very useful" are linearly mapped to 0 to 1 as labels to train a modified implementation of the siamese neural network proposed in Susaiyah et al. [2021]. It uses context means and bag-of-schema words (BoSW) features. The main modification of the network from Susaiyah et al. [2021] to our network is that our network uses context means instead of histograms. Therefore the input dimensions of our model are much smaller and allow insights involving comparison to scalar values. The BoSW is similar to the bag of words Weinberger et al. [2009], but, we pick the context and measurement labels from the schema definition which are more relevant than the surface forms. Neural networks require a lot of training data, hence to adapt to fewer labelled data availability, we follow semi-supervised training Vapnik and Chervonenkis [1974]. Here, we treat the few labelled insights as seed data to turn unlabelled insights into pseudo-labelled insights using a K-nearest neighbours algorithm Goldberger et al. [2004] based on the Euclidean distance of the BoSW features. The pseudo-labelled data is then used to train the network model. In the absence of user feedback, the  $Score_U$  is assigned a value of 1. The overall relevance score of insight is calculated as a product of the completeness, significance and usefulness scores.

$$\begin{aligned}
 Score_C &= \frac{N_{rec}}{F_{exp} * T} \\
 Score_S &= \frac{1}{1 + \exp(-\frac{\gamma\delta}{\tau})} \\
 Score_F &= Score_C * Score_S * Score_U
 \end{aligned} \tag{1}$$

#### 3.3 Surface realisation

This stage corresponds to the how-to-say component of the GENF. This leads to a complete insight text as shown in Figure 1. Here, placeholders like  $\{context : 1\}$ ,  $\{mean : 1\}$ , etc are filled in. The tenses of verbs are modified based on the context. For example, if the context of a verb relates to a period in the past, the past tense is assigned to it. The verbs that have to be modified are mentioned in the template of the form  $\{tense(verb, x)\}$ , where x is 2 if the verb indicates the user (second person) and 3 (third person) if it indicates the measurement.

#### 3.4 Insight recommendation

The insights populated are usually in large numbers, and hence, a selection mechanism is employed in our approach to choosing the best and most diverse set of insights (what-to-say). Here, the insights are clustered based on BoSW features. For clustering, the K-means algorithm Hubert and Arabie [1985] based on Euclidean distance is used. The value of K is set to be equal to the number of insights that we want to present to the user. Finally, we choose to show to the user, the insight having the highest  $Score_F$  from each cluster.

Table 1: Sample set of insights recommended and feedback obtained (FB):👍: useful, 👎: not-useful

Insight	FB
With physician-9 (0.25 hours) the duration of a shift was lower than other physicians (4.27 hours)	👍
On Wednesdays, the duration of an exam in the morning (1.37 hours) is greater than in the afternoon (0.74 hours)	👍
On Thursday mornings (10.03) the acquisition time of an exam was 3.92% lower than on other day mornings (10.44)	👎

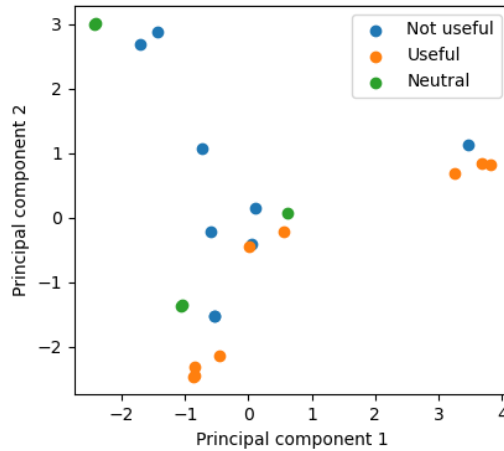


Figure 2: Principal component projections of the insights

## 4 Experiments

### 4.1 Preliminary insights and feedback on usefulness

We used our approach and generated insights on the utilization data of an interventional radiology department of a hospital [Anonymous]. The schemas consider contexts like physicians, exams, periods and patient age over different measurements like exam duration, acquisition time for a scan, radiation dosage and physician shift times. A total of 2470 insight candidates were generated of which 730 were truthful and scored using a  $\gamma$  value of 6.  $Score_U$  is assigned 1 due to the absence of feedback. A 51-element BoSW feature vector is used to select 23 insights by the clustering approach. We then collected feedback from the hospital administration on the usefulness of the insights. A few of these insights and the feedback obtained are shown in Table 1.

The two most dominant eigenvectors Wold et al. [1987] of the insight’s BoSW features are shown in Figure 2. It is observed that all the ‘not useful’ insights seem to be clustered together. This indicates the discriminating power of the BoSW features. The statistics of the feedback are shown in Table 3. It is seen that quite a few insights are not-useful. This can be further improved by incorporating user feedback while recommending the insights. The insights chosen after one round of feedback from the user incorporating the neural network are shown in Table 2. The corresponding statistics are shown in Table 3. Although there aren’t any significant differences in the usefulness. We observed a slight drop in the number of insights that were not useful which could turn significant with multiple rounds of feedback.

Table 2: Insights recommended by our selection approach after one round of feedback from the user

Insight	$Score_U$
With physician-20 (6.55 exams) the number of exams in a shift was greater than other physicians (3.61 exams)	0.67
On Wednesday afternoons the duration of an exam in patients of age 50 to 70 (29.51 minutes) is lower than in patients older than 70 years (51.75 minutes)	0.67
On Wednesday afternoons (9.56 s) the acquisition time of an exam was 13.55% lower than on other day afternoons (11.05 s)	0.8

Table 3: Statistics of the feedback given to the insights

◇: First round of feedback (default recommender)

♡: Second round of feedback (neural recommender)

Statistic	◇	♡
Total Insights	23	23
not at all useful	0	0
not useful	<b>9</b>	8
neutral	5	6
useful	<b>9</b>	<b>9</b>
very useful	0	0

## 5 Discussions

### 5.1 Scalability

Scalability is important for the effective implementation of the OGR paradigm. Schematising the insight generation divides the problem into independent subsets that allow for parallel processing. A comparison of the computation time with and without parallel processing is presented in Table 4. The rate of generating insights by the parallel implementation of GEN-IG is slower than its serial counterpart when the number of insights is less than 700. However, this changes when the insights increase in numbers.

### 5.2 Domain-adaptation

The schema approach introduced here can be extended to other domains by adapting the schema instead of re-coding. Additionally, the generalisation capabilities of large language models could be leveraged to perform this adaptation task in future.

## 6 Conclusion

In this paper, we presented the schema-based approach that allows us to easily define, control and generate actionable insights. We also described a recommendation system that picks diverse, significant and useful insights to be shown to the user using features derived from the schema. In our preliminary experiment, we found a slight improvement in the

Table 4: Performance of parallel vs serial processing

Total insights (count)	Significant insights (count)	Time for serial (min)	Time for parallel (min)
615	176	<b>2.53</b>	2.76
1237	345	3.59	<b>3.12</b>
1855	554	4.65	<b>3.68</b>
2470	730	5.86	<b>4.67</b>

usefulness of the insights with one round of feedback which we believe would improve further with subsequent rounds. In future, we plan to implement this approach in more frequent feedback scenarios such as personal health coaching and campaign management.

## Acknowledgements

This work is supported by the Horizon H2020 Marie Skłodowska-Curie Actions Initial Training Network European Industrial Doctorates project under grant agreement No. 812882 (PhilHumans).

## References

- Rasmus Leth Jørnø and Karsten Gynther. What constitutes an ‘actionable insight’ in learning analytics? *Journal of Learning Analytics*, 5(3):198–221, 2018.
- Aki Härmä and Rim Helaoui. Probabilistic scoring of validated insights for personal health services. In *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1–6. IEEE, 2016.
- Ehud Reiter. An architecture for data-to-text systems. In *Proceedings of the Eleventh European Workshop on Natural Language Generation, ENLG ’07*, pages 97–104, Stroudsburg, PA, USA, 2007. Association for Computational Linguistics. URL <http://dl.acm.org/citation.cfm?id=1610163>.1610180.
- Samuel J O’Malley, Ross T Smith, and Bruce H Thomas. Data mining office behavioural information from simple sensors. In *AUIC*, pages 97–98, 2012.
- Isabel Funke, Rim Helaoui, and Aki Härmä. Interactive health insight miner: an adaptive, semantic-based approach. In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 478–479, 2018.
- Rakesh Agrawal and John C Shafer. Parallel mining of association rules. *IEEE Transactions on knowledge and Data Engineering*, 8(6):962–969, 1996.
- Alex A Freitas. On rule interestingness measures. In *Research and Development in Expert Systems XV*, pages 147–158. Springer, 1999.
- Usama Fayyad, Gregory Piatetsky-Shapiro, and Padhraic Smyth. From data mining to knowledge discovery in databases. *AI magazine*, 17(3):37–37, 1996.
- Allmin Susaiyah, Aki Härmä, Ehud Reiter, and Milan Petković. Iterative neural scoring of validated insight candidates. In *ECAI workshop on Intelligent Information Processing and Natural Language Generation.*, Santiago de Compostela, Spain, Sep 2020a. URL [https://intellang.github.io/papers/6-IntelLanG\\_2020\\_paper\\_6.pdf](https://intellang.github.io/papers/6-IntelLanG_2020_paper_6.pdf).
- Nandan Sudarsanam, Nishanth Kumar, Abhishek Sharma, and Balaraman Ravindran. Rate of change analysis for interestingness measures. *Knowledge and Information Systems*, pages 1–20, 2019.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9, 2019.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Allmin Susaiyah, Aki Härmä, Ehud Reiter, Rim Helaoui, Milan Petković, et al. Towards a generalised framework for behaviour insight mining. In *SmartPHIL: 1st Workshop on Smart Personal Health Interfaces*. ACM, 2020b.
- Albert Gatt and Emiel Kraemer. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research*, 61:65–170, 2018.
- Michael Naaman. On the tight constant in the multivariate dvoretzky–kiefer–wolfowitz inequality. *Statistics & Probability Letters*, 173:109088, 2021.
- Michael P Fay and Michael A Proschan. Wilcoxon-mann-whitney or t-test? on assumptions for hypothesis tests and multiple interpretations of decision rules. *Statistics surveys*, 4:1, 2010.
- The pandas development team. pandas-dev/pandas: Pandas, February 2020. URL <https://doi.org/10.5281/zenodo.3509134>.
- Allmin Susaiyah, Aki Härmä, Ehud Reiter, and Milan Petković. Neural scoring of logical inferences from data using feedback. *International Journal of Interactive Multimedia & Artificial Intelligence*, 6(5), 2021.
- Kilian Weinberger, Anirban Dasgupta, John Langford, Alex Smola, and Josh Attenberg. Feature hashing for large scale multitask learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 1113–1120, 2009.

- 
- Vladimir Vapnik and Alexey Chervonenkis. *Theory of pattern recognition*. Nauka, Moscow, 1974.
- Jacob Goldberger, Geoffrey E Hinton, Sam Roweis, and Russ R Salakhutdinov. Neighbourhood components analysis. *Advances in neural information processing systems*, 17, 2004.
- Lawrence Hubert and Phipps Arabic. Comparing partitions. *Journal of classification*, 2(1):193–218, 1985.
- Svante Wold, Kim Esbensen, and Paul Geladi. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3):37–52, 1987.