

An Ontology to support Knowledge Management in Behaviour-based Healthcare

John Kingston and Nathaniel Charlton

University of Brighton, Brighton, UK
j.k.kingston@brighton.ac.uk

Abstract. A Do Something Different (DSD) behaviour change programme consists of a series of digitally delivered behavioural prompts, or “Dos”, for targeted behaviour change designed to aid people to achieve personal goals or to break bad habits. “Dos” may address for example behaviour connected with exercise, smoking, diet, sleep or diabetes self-management, or personal development objectives such as leadership. DSD’s current database contains thousands of such “Dos”, developed by behaviour change experts.

We have developed an ontology to organise the large database of behavioural prompts in order to help DSD with its knowledge management. We began by developing an ontology of function, bottom-up; then we expanded this into a multi-perspective ontology covering WHO, WHAT, HOW, WHEN, WHERE and WHY perspectives.

The expected benefits from our ontology are: 1. response data from users can be aggregated by ontology category to get insights into users’ behaviour, 2. the ontology enables the implementation of smarter algorithms for selecting which “Dos” to suggest to which users, and 3. the ontology assists in the process of generating new “Dos”, either by domain experts or by crowd-sourcing.

Finally we follow an ontology-theoretic approach to argue that our ontology should be re-usable across other behaviour change applications, but maybe not across other application types.

Keywords: Ontology, Knowledge Management, Healthcare, Health, Behaviour Change, Case Based Reasoning

1 Introduction

Since 2012, Do Something Different Ltd (DSD) has designed and delivered a wide range of behaviour change interventions, addressing health and wellbeing issues such as stress reduction, weight loss and diabetes self-management, and broader personal and organisational development objectives such as leadership, diversity and management skills. It has been argued that the most effective way of changing thoughts and attitudes is via behaviours [1, 2]. Results have been reported in [3] and [4].

Each DSD intervention begins with an online pre-programme diagnostic questionnaire, where the user answers questions about their behaviours, habits, wellbeing, thoughts and feelings. Then over the next few weeks, the participant receives a series

of personalised recommendations of small activities, called “Dos”, that are outside their normal habits [2]. For example, on a smoking cessation programme, a user who has answered that they often smoke “while sitting in your favourite place/chair/spot on the sofa” might be advised “Today break one connection: only smoke when standing up. Don’t take one sitting down!”

“Dos” are sent by smartphone app, email or SMS, and supported with an online forum and motivational messaging. Participants are offered the chance to complete the questionnaire again after their programme; doing so gives them access to a personalised report comparing their pre- and post-programme scores. See Figure 1 for an example.

While many of the “Dos” in a DSD intervention are directly related to the objective of the programme, interventions also aim to promote behavioural flexibility in general. This aspect of the interventions helps people practice behaving in ways they currently do not, or that are outside their comfort zone, such as assertively, proactively or spontaneously. A person who answered that they do not behave assertively might receive a prompt, “*Be a bit more assertive today: Speak up when you would normally hold back. Be direct in asking for what you want.*”

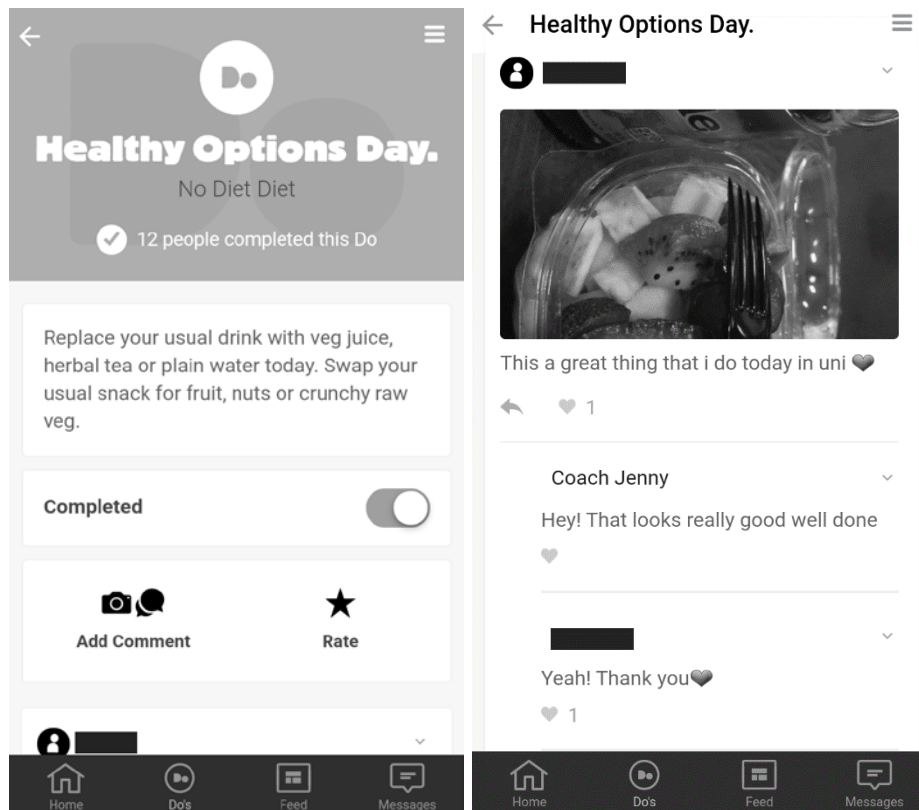


Figure 1: An example Do, plus feedback from user and coach

DSD’s behaviour model was originally based on findings from a series of papers and books by psychologists Fletcher, Pine and others (e.g. [2,5]) which listed 30 behaviours and aimed to promote them all, on the grounds that the more behaviours a person has to choose from, the more behaviourally flexible they were. Subsequent analysis (reported in [4]) found that while most of those 30 behaviours are associated with better wellbeing (in self-report data), not all of them are; thus the DSD model currently promotes twenty-one of the original behaviours (the 21 nearest to the right in Figure 2).




Figure 2: A correlation network showing co-occurrence of the 30 behaviours, from [2]

2 Ontology

Not all the Dos in DSD’s database are explicitly targeted at modifying one of the twenty-one behaviours directly. A first attempt to categorise the Dos produced a list of 130 different categories, such as “Spend time in nature”, “Cut out moaning/complaining” or “Take Risks”. The latter is directly related to one of the 21 behaviours (‘risk-taker’) but the other two are not.

If the majority of the 130 categories of Dos are not directly promoting one of the twenty-one behaviours, then how can they be classified? A further analysis grouped most of the 130 categories into eight top level categories:

1. Actions that directly express one of the 21 behaviours.
2. Actions that disrupt the user’s general routine behaviours.
3. Actions that require physical exercise.
4. Actions that make use of the person’s skills or talents.
5. Actions that express kindness to others.
6. Being social; connecting or sharing with other people.
7. Being aware; noticing things around you that might normally be ignored.
8. Being organised.

Some of these top level categories have subcategories; we were pleased to find that the subcategories predicted by literature or common sense fitted the collection of available Dos. For example, “being kind” can be broken down according to the five typical “love languages” [6]: expressing words of affection; offering someone quality time; giving gifts; performing acts of service; and offering affectionate physical contact. The “being aware” category also breaks down into five obvious groups corresponding to the five senses, although a ‘General’ subcategory was also needed to accommodate Dos such as:

“Tourist Day. Be a tourist in your local town, send postcards, browse souvenirs, take photos, visit at least one attraction you've never been to before.”

It even proved possible to categorise Dos within this scheme that had been difficult to fit into the 130 category analysis, such as:

“Fresh Feet day. Walk barefoot on as many different surfaces as you can today. Feel free to adorn your tootsies too. Record how it looked/felt.” (Be Aware: Touch)

“Do a different walk day. Avoid the cracks in the pavement, or skip for 10 paces in every 50. Walk fast for one minute, slow for the next. Or take a bag and pick up litter.” (Disrupt Routine Behaviour; the litter option adds Be Aware)

“Fix-It Day. Mend at least 1 thing that's broken, or replace spent light bulbs/batteries/refills.” (Use your Skills)

3 Multi-Perspective Ontology

Not all of the 130 subcategories fitted into the above eight top level categories. For example, there is a subcategory that has been loosely named “Smoking related”, consisting of ‘Dos’ designed for people on smoking cessation programmes. Some of them are *disruptor* Dos, provided for people who find themselves desperate for a cigarette and want to be distracted quickly. These Dos are therefore obtainable on demand, reasonably brief to complete, and can be performed immediately. However, the “Smoking related” Dos range across several of the top level categories, particularly those involving exercise, disrupting routines, or being organised.

It is not a problem in principle for Dos to appear in more than one category. However, what is happening here is that Dos are being categorised according to their purpose – to distract people from smoking in the short term – rather than their function (awareness, socialising, exercise, etc.)

Kingston [7] refers to this as a multi-perspective approach to classification. He proposes that a complete classification for any knowledge management application can be obtained by classifying according to five or six different perspective – WHO, WHAT, HOW, WHEN, WHERE and (sometimes) WHY.

If we apply this analysis to the 130 categories of Dos, we observe the following:

- The eight-way classification described in the previous section describes the function of the various Dos. This can be seen as the HOW perspective.
- There are categories of Dos that are deemed relevant only to a single DSD programme – stress reduction, weight loss, smoking cessation or diabetes self-management. Since these are categorised according to the goal that is being aimed

for, which is the purpose of the programme, these can be considered to be the WHY perspective.

- Some categories relate to objects (or groups of objects) in the world. There are food-related Dos; alcohol-related Dos; business-related Dos; environment-related Dos; and digital-related Dos. These constitute the WHAT perspective.

We can also see that some or all Dos could benefit from having some attributes defined:

- The disruptor Dos are designed to be brief and able to be performed immediately. These are two key features of the WHEN perspective: start time (which can be at various levels of detail, from minutes to months) and duration of an activity or event.

Another Do that has a key WHEN attribute is:

“Be more spontaneous today. Do something on the spur of the moment without too much deliberation. Surprise someone by doing something different.”

- The WHO perspective has a similar duality. Some Dos (particularly those related to socialising or kindness) may specify who should be an appropriate target of the Do. Also, some Dos are not appropriate for certain users to perform – strenuous physical exercise should probably be avoided for patients on a weight loss programme, for example.
- Some Dos specify WHERE they should take place, either generally or specifically. For example,

“Be unpredictable today! Wave at a stranger on a passing bus, or out of your window.”

In addition to all the above, it is also possible in theory to record who, what, how, when, where and why the user actually performed the Do, if the user comments on their Do in sufficient detail. Some of this information is automatically available: for example, if a Do is sent out on a Monday and the user reports completing it on a Thursday, the time when the Do was performed is known to an accuracy of 1-2 days.

The thesis of this paper is that applying a multi-perspective ontology to the Dos (i.e. populating six attributes that represent WHO, WHAT, HOW, WHEN, WHERE and WHY) will offer significant benefits for DSD’s management of their behaviour change programmes.

4 Knowledge Management Benefits of a Multi-perspective Approach

A multi-perspective categorisation like this one should enable targeting of Dos to particular users; automating the selection of Dos; and creation of further Dos.

4.1 Targeting Dos to particular users

The Holy Grail for an organization offering behaviour change programmes is to select interventions that will have the maximum effect for a particular user to achieve a particular goal. Having a multi-perspective ontology can help significantly with such selections. The reason is that the various attributes of Dos can be mapped against results of the behaviour change programmes and also between themselves, making it possible to determine if (for example) exercise-related Dos have more effect in the summer months, or social-related Dos are more effective if they take place outside work rather than inside work.

4.2 Automating selection of Dos

DSD is currently testing a change where users are offered a choice of three Dos instead of being sent just one. The rationale behind this change is that certain users like some Dos better than others.

Using a multi-perspective ontology, it would be possible to ensure that the three Dos offered are significantly different from each other. This would be done by making sure the three Dos all had different values on each one of the different attributes.

Another approach would be to combine automated selection with targeting of Dos by supplying Dos that are similar to ones that the user has previously liked. It is important to select Dos that the user is going to enjoy, as assuring engagement with the programme is one of the biggest hurdles that DSD faces. This could be done using a simple case based reasoning approach.

This approach would appear even more intelligent if an extra layer of dimensionality was given to the attributes. For example, all attributes could be given dimensions of “safe” through to “extreme”. For exercise, walking would be “safe”, sprinting or wrestling would be “extreme”. For kindness-related Dos, physical affection might be considered “extreme”. For socializing Dos, casual contacts would be “safe” while making a long-term commitment to a group that required significant amounts of time (a drama production, perhaps) might be “extreme”.

If the user selects “extreme” Dos two or three times in a row and/or indicates that s/he enjoyed them, then it would be a good idea to include in the next choice of three a Do that is “extreme” on a different attribute.

4.3 Creating new Dos

Finding the inspiration to expand an already extensive set of Dos is a difficult task. After a while, duplication of existing Dos is a significant risk; indeed, DSD’s database already contains some Dos that are virtually duplicates of others.

However, with a multi-perspective categorization, it is possible to list Dos according to all their categories. If some categories have few or no Dos, it is a much easier task to think of new Dos for those categories, because the categories themselves act as priming information. For example, one category would be: being kind to someone by spending quality time with them – related to food – at work – WHO and WHEN are left to the user’s discretion. This might lead to a new Do of “*Buy lunch for someone at work today who you don’t normally talk to. Eat your lunch with them and do your best to listen*”

more than you talk.” The WHY perspective would be the goal of the programme; this Do might not be appropriate for a diabetes self-management programme.

Another method to create new Dos is crowd-sourcing i.e. allowing users who have experience of the programme to create their own Dos by simply deciding that they will do something different, of their own choosing. Users are encouraged to feed back their new Dos to the development team, who can decide whether to add them to their database of Dos.

Deciding whether the new Do is truly new or is a duplication or instance of an existing Do is a large task when there are 2500 Dos. However, if the new Do can be classified according to the multi-perspective ontology, the number of existing Dos to compare for duplication falls to a much more manageable number.

In theory, a language processing application might be developed to help in automatic classification of these potential new Dos. In practice, the volume of new Dos is expected to be no higher than 10 per week, and the effort required to add six attributes to each of 10 or fewer Dos is insufficient to justify the development of intelligent supporting software.

5 A Perspective from Ontology theory

The development of ontologies has frequently been proposed as an aid to knowledge management but much less frequently implemented. The reasons are usually to do with issues that relate to the theory of ontologies.

5.1 Reusability versus Specificity

Ontologies were once touted as the next big thing in knowledge management [8] [9] because they were supposed to be re-usable. The idea was that a knowledge categorization would be defined once and could then be re-used by multiple applications.

It turned out that this did not work well. Categorisations were often domain-specific. Sometimes this was because there was no underlying real-world categorisation, and so arbitrary categories had to be invented; Wielinga [10] gave an example from the Rijksmuseum in Amsterdam of experts in paintings using a completely different classification of colour from experts in ceramics. Colour does not have fundamental categories in the real world; rather, it is formed by the perception of three wavelengths of light. So only a handful of colour categories can be unambiguously defined (the eight categories that involve all or none of each wavelength); other categories must be agreed.

There were other reasons too. Sometimes the parties creating the classification had a vested interest in seeing their area of interest classified as high up the tree as possible [11]. Sometimes the issue was the lack of a multi-perspective classification, particularly relating to purposes, leading one ontology to classify a strut (within a building) as a “beam” while another might classify it as a “support” [12].

When we apply these insights to our multi-perspective ontology based on behaviour change, new opportunities can be identified. Our ontology is really an ontology of activities that happen to be directed towards goals in healthcare and wellbeing; so the attributes and (perhaps) range of values for at least two of the six perspectives (WHEN

and WHERE) can be re-used directly from other activity-based ontologies. Reuse may be possible for the WHO perspective as well.

The perspectives that are directly related to healthcare are HOW (the eight categories identified in section 2) and WHY (the goal of each programme). The WHY perspective is uncontroversial; it is an (uncategorized) list of health benefits that are commonly sought by large groups of people. So the key ontological benefit for healthcare from this project are the eight categories of the HOW perspective.

Are these re-usable? It seems likely that they will be re-usable in other behaviour change programmes. It is unclear whether it might be of benefit in programmes such as cognitive therapy, which focus on modifying thinking and beliefs rather than on changing behaviours. We claim that our ontology is likely to be re-usable for similar types of application but make no further commitment.

5.2 Mixing of perspectives within classifications

A second problem is that many ontologies mix up taxonomies, mereonomies and other relation-based structures. [13] reviewed the ACM Digital Classification and showed how the subcategories mixed categories based on programming techniques; categories based on applications of programming; and, in the case of the Hardware category, subcategories based on parts.

The ontology presented in this paper largely avoids mixed perspectives; its breakdowns are heavily based on methods for achieving the top level goal (i.e. they all follow the HOW perspective). The ‘five love languages’; the five senses (for Awareness interventions); different types of exercise; different ways of using skills and talents; different ways to be organized; and different ways of disrupting routines are all method-based breakdowns. This increases the possibility that the different categories can each support knowledge management in the same way, making the ontology easier for DSD to apply. It also suggests that overarching dimensions (such as “safe” versus “extreme”) might be developed that apply across all categories.

5.3 Level of abstraction

A third problem that arises with ontologies is that the level of interest that an application took in a particular category affected how it was classified within the ontology. To quote [14], a car salesman may view ‘colour’ as merely an attribute of the concept ‘car’, while a photographer might view ‘colour’ as a concept in its own right with attributes including ‘hue’ and ‘saturation’.

When something is a concept in one ontology but a property in another, this implies that the two ontologies are one level of abstraction apart regarding this thing. Guarino [15] identifies no fewer than nine levels at which an ontology may exist, ranging from the atomic and static levels (where ‘wavelengths of light’ belong) right up to intentional and social levels (at the latter of which social rules and conventions apply). These are shown in Figure 3.

The Dos can all be said to exist at the intentional level since they require intentional behaviour, although few refer directly to social rules or conventions. However,

some social rules need to be applied on a day when a user is expected to choose someone to whom to show physical affection. So our ontology is at least internally consistent in terms of abstraction.

Is our ontology re-usable? The answer is that it should be re-usable for other behaviour change programmes because such programmes (by definition) require intentional behaviour. However, it is probably not that helpful for an application at a much lower level of abstraction, such as a natural language analysis application that determines the parts of speech present in a typical ‘Do’ or other instructed activity.

| |
|---|
| Atomic |
| Individuation: minimal size |
| Persistence: spatio-temporal continuity |
| Static |
| Individuation: mereological sum of atoms |
| Persistence: same properties |
| Mereological |
| Individuation: mereological sum of entities |
| Persistence: same parts |
| Topological |
| Individuation: self-connection |
| Persistence: similar topology |
| Morphological |
| Individuation: proximity |
| Persistence: similar shape |
| Functional |
| Individuation: purpose |
| Persistence: persistence of functionality |
| Biological |
| Individuation: presence of life |
| Persistence: persistence of life |
| Intentional |
| Individuation: intentional behavior |
| Persistence: persistence of intentional behavior |
| Social |
| Individuation: inter-agent connections |
| Persistence: persistence of inter-agent connections |

Figure 3: Guarino’s nine ontological levels

6 Conclusions and future work

We have shown how a multi-perspective ontology can provide significant knowledge management benefits to an organisation that offers healthcare through behaviour change. It can help generate new interventions and select interventions that are

appropriate for a particular user to achieve a particular goal. We believe that it can also help improve engagement with the programme.

We believe that our ontology is re-usable across other behaviour change applications but we will not commit to it being useful for other types of application.

In future work we hope to report on interaction and outcome data from newly revised programmes, making use of our ontology, to test our beliefs empirically.

Acknowledgements

The authors were supported by Innovate UK via Knowledge Transfer Partnership 10152.

Thanks are due to Professor Ben (C) Fletcher for comments on this paper.

References

1. Pine, K., Fletcher, B. Shifting brain channels to change health behaviour. *Perspectives in Public Health*, 134, 1, 16-17 (2014). DOI: 10.1177/1757913913514705
2. Fletcher, B., Hanson, J., Pine, K., Page, N: FIT – Do Something Different: A new psychological intervention tool for facilitating weight loss. *Swiss Journal of Psychology* 70(1), 25–34 (2011).
3. Pine, K., Fletcher, B., Changing people’s habits is associated with reductions in stress, anxiety and depression levels. Technical Report, Do Something Different Ltd (2016)
4. Charlton, N., Kingston, J., Petridis, M., Fletcher, B.: Using Data Mining to Refine Digital Behaviour Change Interventions. In *Proceedings of 7th International conference on Digital Health 2017*, pp. 90 – 98, ACM, LONDON, United Kingdom (2017).
5. Fletcher, B., Pine, K.: *Flex: Do Something Different*. University of Hertfordshire Press, Hertfordshire, (1999).
6. Chapman G.: *The Five Love Languages: How to Express Heartfelt Commitment to Your Mate*. Northfield (1995)
7. Kingston, J.: *Multi-Perspective Modelling for Knowledge Management and Knowledge Engineering: Practical Applications of Artificial Intelligence*. CreateSpace (2017).
8. Brandt, S., Morbach, J., Miatidis, M., Marquardt, W.: An ontology-based approach to knowledge management in design processes. *Computers & Chemical Engineering* 32(1), 320-342 (2008).
9. Abecker, A., van Elst L., *Ontologies for Knowledge Management*. In: Staab, S., Studer, R. (eds.) *Handbook on Ontologies*. International Handbooks on Information Systems. Springer, Berlin, Heidelberg (2004).
10. Wielinga R., University of Amsterdam. Personal communication to the first author of this paper. (1997)
11. Kingston, J., *Merging top level ontologies for scientific knowledge management*. In: *Proceedings of the AAAI Workshop on Ontologies and the Semantic Web*, pp. 147-154 (2002).
12. Kingston, J., *Multi-perspective ontologies: Resolving common ontology development problems*. *Expert Systems with Applications* 34 (1), 541-550 (2008).

13. Kingston, J., Ontology, knowledge management, knowledge engineering and the ACM classification scheme. In: Research and Development in Intelligent Systems XIX, pp. 207 - 220 (2003).
14. Kingston, J. Multi-Perspective Modelling for Knowledge Management and Knowledge Engineering: Practical Applications of Artificial Intelligence. CreateSpace, 2016. ISBN 978-1539048343.
15. Guarino, N., Some Ontological Principles for Designing Upper Level Lexical Resources. In: Rubio, Gallardo, Castro, Tejada (eds.), Proceedings of First International Conference on Language Resources and Evaluation, pp. 527-534, Granada, Spain (1998).