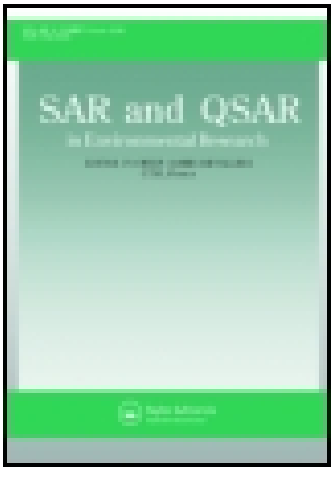

SAR and QSAR in Environmental Research

Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/gsar20

ToxRead: A tool to assist in read across and its use to assess mutagenicity of chemicals

G. Gini[a], A.M. Franchi[a], A. Manganaro[b], A. Golbamaki[b] & E. Benfenati[b]

[a] Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Milan, Italy

[b] Dipartimento di Salute e Ambiente, IRCCS - Istituto di Ricerche Farmacologiche 'Mario Negri', Milan, Italy

Published online: 15 Dec 2014.

To cite this article: G. Gini, A.M. Franchi, A. Manganaro, A. Golbamaki & E. Benfenati (2014): ToxRead: A tool to assist in read across and its use to assess mutagenicity of chemicals, SAR and QSAR in Environmental Research, DOI: 10.1080/1062936X.2014.976267

To link to this article: http://dx.doi.org/10.1080/1062936X.2014.976267

SAR and QSAR in Environmental Research, 2014
http://dx.doi.org/10.1080/1062936X.2014.976267

ToxRead: A tool to assist in read across and its use to assess mutagenicity of chemicals[£]

G. Gini[a], A.M. Franchi[a], A. Manganaro[b], A. Golbamaki[b] and E. Benfenati[b]*

[a]*Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Milan, Italy;*
[b]*Dipartimento di Salute e Ambiente, IRCCS – Istituto di Ricerche Farmacologiche 'Mario Negri', Milan, Italy*

(*Received 5 August 2014; in final form 15 September 2014*)

Life sciences, and toxicology in particular, are heavily impacted by the development of methods for data collection and data analysis; they are moving from an analytical approach to a modelling approach. The scarce availability of experimental data is a known bottleneck in assessing the properties of new chemicals. Even when a model is available, the resulting predictions have to be assessed by close scrutiny of the chemicals and the biological properties of the compounds concerned. To avoid unnecessary testing, a read across strategy is often suggested and used. In this paper we discuss how to improve and standardize read across activity using ad hoc visualization and data search methods which use similarity measures and fragment search to organize in a chart a picture of all the relevant information that the expert needs to make an assessment. We show in particular how to apply our system to the case of mutagenicity.

Keywords: read across; mutagenicity; REACH; similarity; VEGA; SAR; QSAR

1. Introduction

According to recent legislation, such as the Registration, Evaluation, Authorisation & restriction of CHemicals (REACH) regulation of the European Union, *in silico* methods are allowed in various situations to derive the properties of new chemicals from all available knowledge. *In silico* methods are different in their role and use [1,2]. They do not require experimental testing in the laboratory, but exploit knowledge obtained from existing *in vivo* or *in vitro* tests.

The most used *in silico* methods are structure–activity relationship (SAR) and quantitative structure–activity relationship (QSAR) which build a model of the phenomenon and give a result in terms of dose or toxicity class. There is a continuous search for good models, for instance, to perform prioritization or screening of large data sets [3]. Good models should:

(1) explain patterns in data;
(2) correctly predict the results of new experiments or observations;
(3) be consistent with previous theories.

*Corresponding author. Email: benfenati@marionegri.it
[£]Presented at the 16th International Workshop on Quantitative Structure-Activity Relationships in Environmental and Health Sciences (QSAR2014), 16–20 June 2014, Milan, Italy.

Certainly the third requirement is a critical one. When research is guided by broad general principles, as in many areas of science and engineering, the models often embody these principles. In several biological and toxicological processes the theory behind the biological interactions is not clearly determined. Therefore, models derived from mathematical techniques have been proposed without cause–effect relationships. These models are, however, subject to criticism of the lack of 'deep knowledge' among scientists.

In order to make an assumption on the output of a model, the experts usually take similar cases and reason on the possible result of the required property as a weighted average of them; this method is a kind of read across. This is done after molecules of interest for the case of study have been collected and organized on the basis of both their chemical structure and their known activity; this is called profiling and it is effective only when supported by a good database. The expert attempts to identify the most similar cases with respect to chemical structure, presence of functional groups, applicability of specific alerts, reasons for considering the parent compounds or its metabolites, and many more. This process is time-consuming and not easy to replicate.

To improve this, some automatic systems have been developed. For instance, the Organisation for Economic Co-operation and Development (OECD) QSAR toolbox provides different ways for profiling; however, they all require the user to interact intensively [4]. Other tools, such as Ambit, provide specific methods to help the expert to accept or rule out the answer of predictive models [5]. However, Ambit is not specifically designed for read across but for assessing the applicability domain of a QSAR model.

To describe the main goal of the presented work, we start from the result of a recent workshop [6] which evaluated the read across method used in the REACH registration process. This workshop concluded that several problems are common with QSAR methods, in particular the applicability domain definition and the categorization. For this purpose, it is useful to refer to structural elements such as functional groups, to evaluate if there are common biomodifications, and to consider the regularity in the way the property changes in the category. Therefore, a well-defined measure of similarity should take into account not only the structure but also common chemical reactions and exclusion rules.

The workshop indicated at least two problems encountered when filling data gaps with read across. The first one is the difficulty in assessing the absence of toxicity, which seems to require a greater burden of proof for justification; this drawback indeed is not a problem of the read across method, but refers to regulators who are more sceptical about accepting a negative result. The second one is how to deal with uncertainty and to what extent results are to be considered reliable. Different elements contribute to reliability: the quality and number of the experimental data used to perform read across; the chemical similarity measures used; knowledge about how chemicals interact with biological systems; and supplementary data from other properties or *in vitro* assays. Naturally, this information is not always available, but each element may contribute in a weight of evidence approach. Uncertainty should be considered in a more systematic way even when read across is used.

Considering the above issues we have developed a system called ToxRead. This aims to be:

- an easy way to obtain and integrate the available knowledge;
- a systematic way to indicate the uncertainty of the result;
- a reproducible way to categorize the substances.

This paper discusses the problem of assessing the mutagenicity property of chemical compounds. This can be achieved by a quick overview of the selected molecule in a space defined by similar molecules.

There is a common accepted theory expressing that the presence of specific functional substructures determines the toxicity of a compound; literature on this abounds in the case of mutagenicity. However, often such structural alerts overestimate mutagenicity, because typically the percentage of experimentally mutagenic compounds having a given structural alert is not 100%, and for some alerts, many of the chemicals possessing the feature are not mutagenic [7]. In these cases it is important to examine both the mutagenic and non-mutagenic cases to make a judgement. Our tool, ToxRead aims to assist the user in making this judgment.

2. Assessing mutagenicity

Human experts usually estimate toxicity on the basis of the detection of structural fragments already known to be responsible for the toxic property under investigation. Such fragments are referred to as structural alerts [8], toxicophores [9] or biophores [10]. These can be determined by human experts from knowledge of the biochemical mechanism of action, such as the activation of an enzyme cascade or the opening of an ion channel.

Mutagenicity, carcinogenicity and reproductive toxicity are some of the most important endpoints to evaluate toxicity towards humans; indeed, they are part of the CMR (Carcinogenic, Mutagenic, Reprotox) regulatory assessment. In particular, mutagenic toxicity is the ability of a substance to cause genetic mutations and it is of considerable public concern due to its close relationship with carcinogenicity [11,12]. An important application is drug and pesticide discovery where the development of candidate compounds, which are potentially mutagens or carcinogens, should be detected as soon as possible during the process.

Mutagenic toxicity can be experimentally assessed by various test systems; the most common is the Ames test [13], which makes use of a genetically engineered *Salmonella typhimurium* and *E. coli* bacterial strains. This test is an *in vitro* model of chemical mutagenicity and consists of a range of bacterial strains sensitive to a large array of DNA-damaging agents [14]. The estimated inter-laboratory reproducibility of Ames test data is 85% [15].

Besides *in vitro* testing, there are also *in silico* methods for mutagenicity such as QSAR and SAR. These make use of a wide variety of statistical methods and a large range of molecular descriptors. One of the first attempts to model mutagenicity for nitro-aromatic compounds [16] used only four descriptors, namely the energy level of the lowest unoccupied molecular orbital (LUMO), the partition coefficient between octanol and water (log P), a structural indicator and a descriptor to exclude molecules considered outliers.

The SAR approach involves the discovery of particular structural fragments in molecules already known to be responsible for the toxic property under investigation. In the mutagenicity/carcinogenicity domain, the most significant innovation in the definition of such toxicophores has been the use of both the formal link between the chemical and toxicological processes, and the extension of the list of moieties covered [14]; this innovative view has been extended by following studies and has inspired many researchers. Starting from the electrophilicity theory of chemical carcinogenesis [17], which correlates the presence of electrophiles (like halogenated aliphatic or aromatic nitro substructures) to genotoxic carcinogenicity, Ashby and Tennant compiled a list of structural alerts for DNA reactivity [18]. In a this study, a few hundred compounds from the US National Toxicology Program (NTP) were mined manually to confirm

the role of structural alerts in the mutagenicity processes; however, the authors did not present numerical correlations between individual substructures and mutagenicity.

Every subsequent effort has started from the knowledge collected by Ashby to derive more specific rules. A more recent study combined an understanding of the mechanism of mutagenesis with statistical criteria [9]. The data set includes more than 4000 molecules with the respective Ames test binary results. A drawback is that molecules tested with different methods (with and without metabolic activation with S9) were mixed; however, it is widely accepted and used in the scientific community. From these core data, several other papers were published which introduced different data mining techniques and extended the collections of rules [19–21].

If the aim is to use mutagenicity as an indicator of carcinogenic substances, the correlation between mutagenicity and rat carcinogenicity is minimal [22]. However, most of the structural alerts are the same on both endpoints.

Practically, structural alerts are rules which state the condition of mutagenicity depending on the presence of a specific chemical substructure. The mutagenicity structural alerts are hypotheses derived from chemical properties and have some degree of mechanistic interpretation. Nevertheless, their presence itself does not give a definitive method to prove the mutagenicity of a compound towards bacterial cells, since other substituents may change the classification. For instance, Snyder et al. [23] reported the results of checking the main commercial rule-based systems for predicting the mutagenicity of pharmaceutical compounds and found that the sensitivity of all systems was low. Moreover, in many cases compounds were found to be mutagenic in the absence of structural alerts. Within another exercise, a number of programs have compared predicting a set of chemicals composed of more than 6000 compounds [24] and, more recently, using the chemicals registered for REACH [25]. The results obtained with some models, like VEGA, have reported accuracy very similar to the reproducibility of the Ames test, indicating that it is possible to predict mutagenicity in many cases [25].

ToxRead is aimed at providing evidence on the evaluation of the relevance of the different structural alerts for the specific chemical of interest, indicating at the same time the most similar compounds which contain these structural alerts.

3. Principles and basis of the proposed software

The developed tool is based on an application supported by libraries of fragments which visualize the substances and the structural alerts.

3.1 *Structural alert libraries*

The ToxRead software is designed to manage multiple libraries. Currently the program includes the mutagenicity libraries present in the VEGA QSAR software [26,27]. These are the libraries taken from the Toxtree software (version 1.60), with the addition of Benigni–Bossa [7] and SARpy rules [28]. The former are derived from human experts and originated from the Ashby set of rules, both extending their number and introducing exclusion rules for some structural alerts.

Given a training set of molecular structures, with their experimental activity binary labels, SARpy generates every substructure in the set and mines correlations between the incidence

of a particular molecular substructure and the activity of the molecules that contain it. This is done in three steps starting from the structural SMILES notation:

(1) Fragmentation. This recursive algorithm considers every combination of bond breakages working directly on the SMILES string. This fast procedure is capable of computing every substructure of the molecular input set;
(2) Evaluation. Each substructure is validated as potential structural alert on the training set. It performs a complete match against the training structures for assessing the predictive power of each fragment;
(3) Rule set extraction. From the huge set of substructures collected, a reduced set of rules is extracted in the form: 'IF contains <structural alert> THEN <apply activity label>'.

The main difference between the Benigni–Bossa and the SARpy rules is that SARpy also generates rules associated with non-mutagenicity. These are conceptually similar to the exclusion rules present in the Benigni-Bossa rules, but the exclusion rules within Toxtree are always associated with a positive toxic rule, while the rules for 'non-toxicity' listed by SARpy are more general and apply to all chemicals.

3.2 *Database of experimental values*

ToxRead makes available the most similar compounds with a certain structural alert. Currently, the experimental values are those extracted from the ANTARES project [29], which refers to the data from Hansen et al. [30], checked and pruned for the chemical structures.

3.3 *Implementation*

The ToxRead tool has been developed as a Java standalone application. This programming language has been chosen because it provides the possibility of running the application on different operating systems without deploying different versions of the application. Furthermore, some libraries of particular relevance were already available in Java.

The application relies on the VEGA core library, which already implements the similarity index. Additionally, the library provides useful features for chemo-informatics purposes such as parsing of SMILES string, SMARTS matching and molecule depiction. This VEGA library itself relies on the CDK (Chemistry Development Kit) libraries [31]. The application uses also the JUNG (Java Universal Network/Graph) framework [32] for the creation and visualization of the chart. The database containing the available compounds and their experimental values was implemented as a local database with the HyperSQL [33] libraries.

In the first beta version of the tool, the database contained 6062 molecules together with the experimental data for mutagenicity. In addition other experimental data have been added: octanol–water partition coefficient (available for 1384 molecules); bioconcentration factor in fish (for 373 molecules); and carcinogenicity (for 531 molecules). This database is stored as a HyperSQL database file of about 16 MB. Access to the database has acceptable performance; for instance, less than 5 seconds are required on a PC with an i3 3.30 GHz core and 8 GB RAM for the main calculation. This consists of accessing the database, calculating the similarity of the target molecule with the available compounds in the database, and extracting the needed molecules to build the chart. The reason why we chose to use a local database is to avoid information exchange through the Internet; we found from our experience that users often work with highly confidential data and are not willing to send any information on external servers for elaboration.

4. Use of ToxRead

4.1 *General features*

ToxRead has been designed to be user-friendly. Colors and shapes represent easily understandable information, with the meaning explained below.

The target chemical is drawn in the center of the visualization panel; it is represented by a circle (see the example given in Figure 1), with several outgoing links to various similar chemicals. The user can choose the maximum number of similar compounds up to six; in fact, from our experience usually three similar chemicals are sufficient. These chemicals are identified using the algorithm implemented in VEGA and the similarity value is calculated as the weighted combination of a fingerprint, three structural keys based on molecular descriptors, and a series of other descriptors (constitutional, hetero-atoms and specific functional groups considering the number of some features or functional groups and not only their presence/absence). The description of the similarity algorithms is presented by Floris et al. [34].

The size of the circle is proportional to the similarity in order to make the user aware of the relevance of each chemical. Clicking on a chemical, the user can see its structure, the similarity value and the experimental values associated with it. The color of the circle indicates whether the chemical is mutagenic (red) or not (green). This color-coding refers only to the experimental value contained in the internal database; the circle is split into green and red when the result is equivocal. Moreover, all the available experimental values such as carcinogenicity are shown and more robust evaluations can be accomplished by the user.

However, these *N* similar chemicals are not the core of the ToxRead tool; indeed the same information is already available within the VEGA software. These *N* chemicals are identified on the basis of a generic, similarity function, without any use of a structural alert.

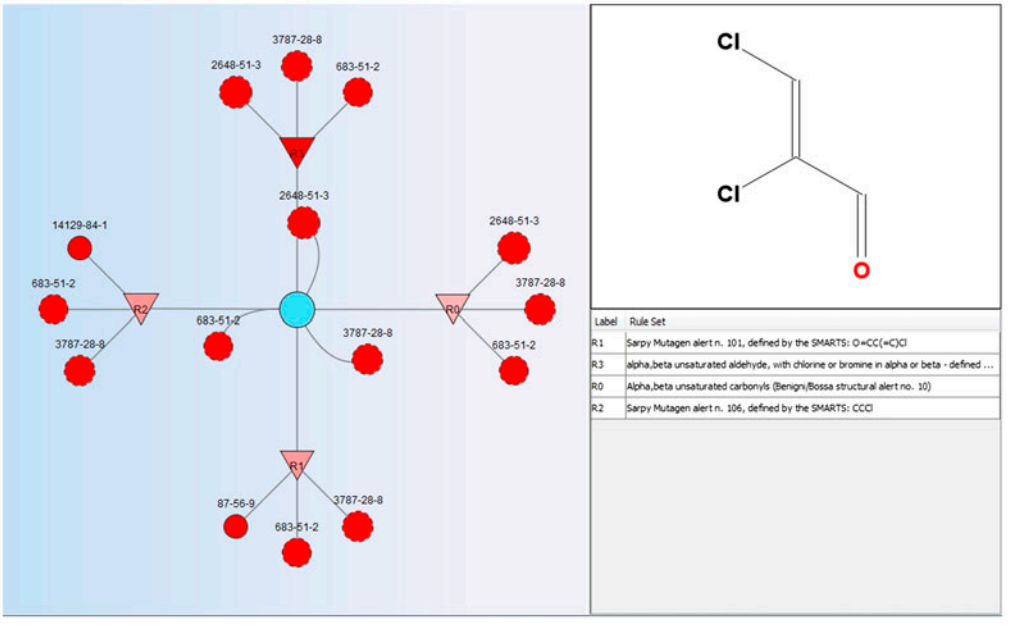

Figure 1. ToxRead screen showing the similar compounds and the rules found in the analysis of the molecule O=CC=C(Cl)Cl.

The innovative idea of ToxRead is that the target chemical is also linked to several structural alerts, giving the user significant additional information. The structural alerts linked to the target chemical are represented by triangles; those pointing upward are non-mutagenic and those pointing downward are mutagenic. In addition, mutagenic alerts are red while non-mutagenic are green. It is important to note that non-mutagenic alerts are not just the absence of mutagenic alerts; we discuss this concept in the following examples.

Another immediate visual detail about the 'validity' of a certain alert is that the saturation of the color is proportional to the percentage of toxic or non-toxic chemicals. It may occur that a mutagenic structural alert refers to a number of chemicals which are mainly non-toxic, and there are examples among the Benigni–Bossa rules that we will see below, but the color of the structural alert is still red because this rule is formally a rule of toxicity; in this way the user is informed about the presence of false positives. This does not apply to SARpy rules, since all of them are by definition positive only if there is a prevalence of toxic chemicals.

By clicking on a structural alert, the user can visualize its chemical structure, its explanation, and the *p*-value relative to the toxicity; by clicking on a specific button it is also possible to visualize up to 100 similar chemicals presenting that structural alert.

ToxRead takes into account the fact that the same similar chemical may appear more than once, linked to different rules. For instance, a nitro-aniline compound will be linked to both the structural alerts of aniline and nitro-aromatic compounds; in these cases the circles are drawn with dashed line. Thus, the overall evaluation is easy when the information presented by the structural alert and by the similar compounds is the same; otherwise users should apply their knowledge to read the ToxRead results and to take the final decision.

In this last case of conflicting results, we expect the user to be conservative, giving emphasis to toxicity alerts. Although we see below how this may not always be the case and the software may help gain a more realistic evaluation.

4.2 *Examples of assessing the mutagenicity property*

Figure 1 illustrates the ToxRead screen for assessing the mutagenicity of the target compound, whose structure is illustrated in the right-hand panel. The figure shows a graph with the three most similar molecules represented as circles, linked to the target compound, which is in the center; the four rules are represented by triangles, each one connected to three other circles. The CAS number of each chemical is shown next to its circle. On the right side of the figure, just below the structure of the target chemical, the list of the associated rules is indicated.

In this case all rules are related to toxic effects. Though, rule R0 is quite generic with a lower percentage of mutagenic compounds (49%; thus most of the chemicals with this fragment are non-mutagenic), while rule R3 is more saturated because it is associated with a higher number of mutagenic compounds. Indeed, rule R3 has an accuracy equal to 1, meaning that all chemicals containing this fragment are mutagenic; looking at the structure of this rule, we can see that it is more specific for the chemical under evaluation, as it contains chlorine or bromine linked to the double bond. In this example all the similar chemicals linked to rules R0 and R3 are those already identified simply considering the similarity according to VEGA. The overall evaluation is quite simple. There are four rules all indicating mutagenicity and all the similar chemicals are mutagenic; the experimental data on the target compound indicates that it is a mutagen.

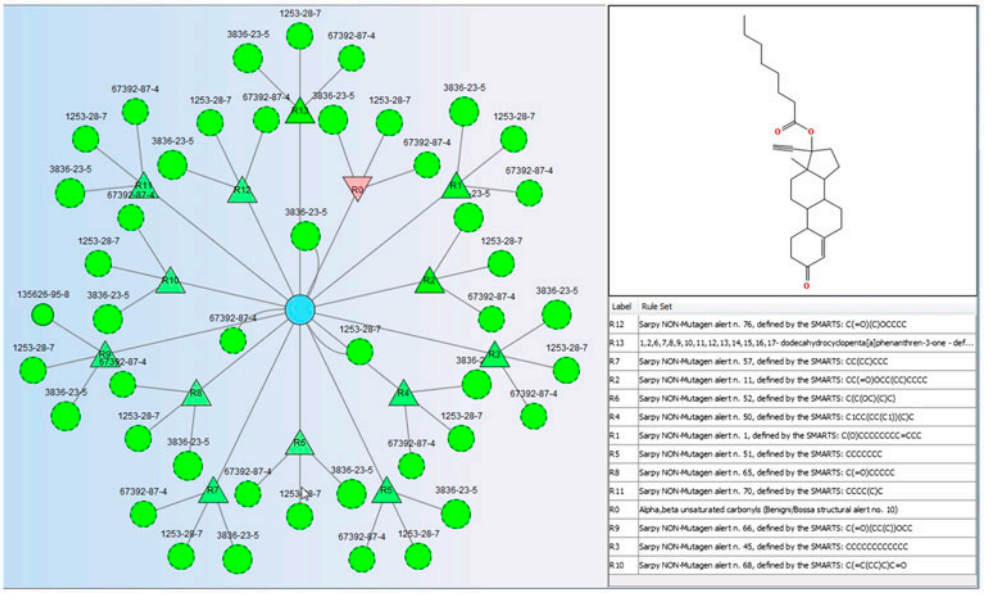

Figure 2. ToxRead screen for assessing the mutagenicity property of the molecule shown on the right.

The second example is given in Figure 2; in this case there are many more rules, with all but rule R0 indicating non-mutagenicity. Observing the graph, however, rule R0 has a high number of false positives, so the user can proceed beyond the generic statistics of this rule and look at the specific similar chemicals containing the structural alert of interest (the unsaturated oxo group). Figure 3 shows these three chemicals linked to rule R0: they are all non-mutagenic. Considering the other similar compounds in Figure 2, it is evident that all of them are non-mutagenic and so there is a broad consensus on the fact that all are non-mutagenic.

In Figure 3 the only critical issue is related to rule R0, but the user can easily rule out this issue in this particular case because rule R0 is generic, and in our case, refers to non-mutagenic chemicals even if they present the mutagenic rule R0. Thus, the user is informed of the possible presence of rules of concern and can decide on the basis of a weight of evidence approach. The elements for the evaluation are the combination of the rules and of the mutagenicity status of the similar chemicals.

The third example (Figure 4) depicts a more complex case. The target chemical is associated with two conflicting rules: the first is the Benigni–Bossa of the aromatic amines and the second derives from the evidence of a group of chemicals sharing the heteroaromatic bicyclic structure. Rule R1 is shown in Figure 5.

This last rule has been defined by us within an ongoing exercise that aims to derive a large set of rules by manual evaluation of chemical classes (in particular aromatic amines) and sub-classes which do not belong to the majority of the class. This approach is the same adopted by Benigni and Bossa for the definition of the exception rules. Figure 5 shows the most similar chemicals found by the software with respect to rule R1; one of them, CAS 703-83-3, is almost the same as the target chemical with the only difference in the methyl group linked to the aromatic ring. In this case ToxRead identifies a toxicity rule, R0, and an exclusion rule, R1, which may be sufficient to overrule the mutagenicity one. Moreover, the

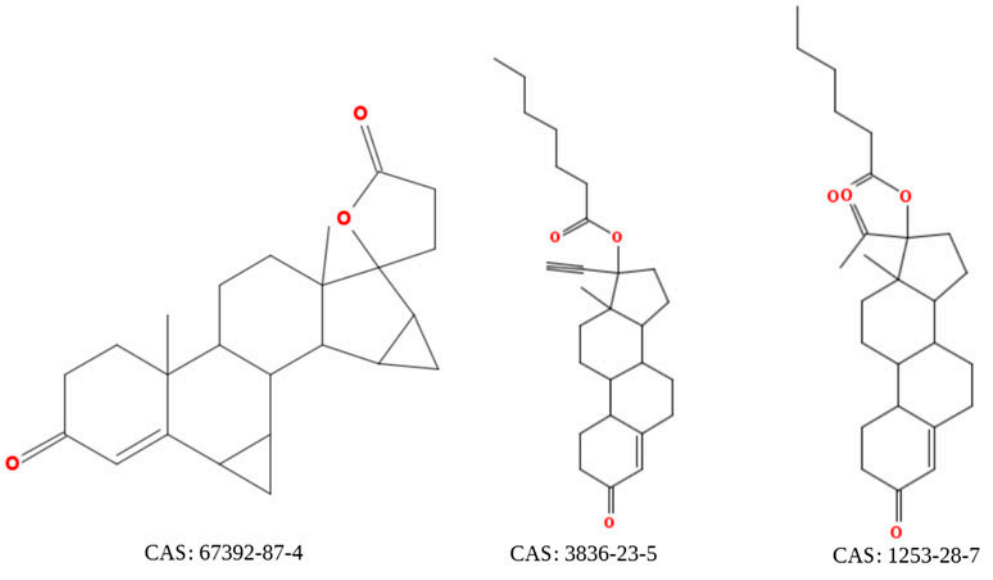

Figure 3. The three most similar chemicals to the target compound containing rule R0, referring to Figure 2.

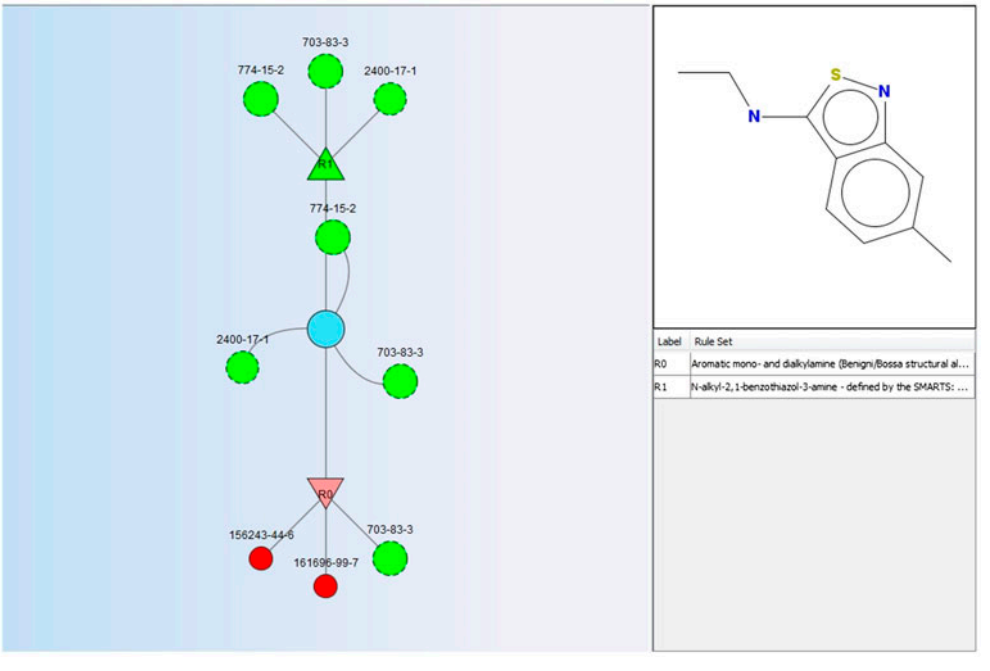

Figure 4. ToxRead applied to a controversial chemical.

evaluation of the three most similar chemicals to the target compound shows that they are all non-mutagenic. These exercises suggest that ToxRead can assist the user in the evaluation of

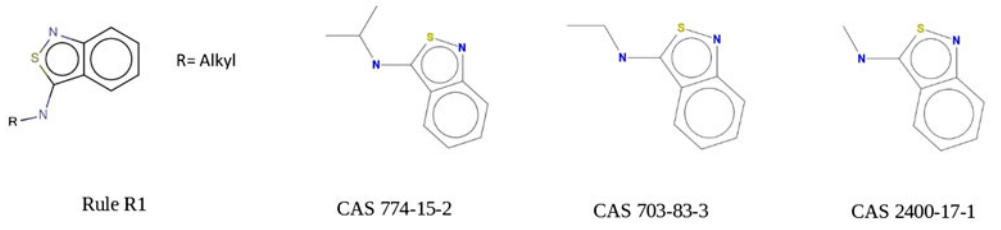

Figure 5. The structure of rule R1 and the three most similar chemicals to the target compound, referring to Figure 4.

conflicting rules for read across, providing a simplified visualization of different pieces of information.

In order to have a sort of validation of our work, we have also applied the mutagenicity models of VEGA to the three molecules used in the above examples. VEGA has three models for mutagenicity: CAESAR, SARpy and Toxtree.

For the first chemical all three models gave the same result as we had obtained from ToxRead; moreover, looking at the most similar chemicals shown by VEGA there was consensus about the mutagenicity of the target compound. The QSAR models and the read across approach supported each other.

For the second chemical both CAESAR and SARpy agreed that the chemical was not mutagenic and the applicability domain index (ADI) was high. Conversely, Toxtree labeled the chemical as mutagenic, but with a very low ADI. The similar compounds were experimentally not mutagenic and thus the overall assessment based on the QSAR models was that the chemical was not mutagenic, but the degree of reliability of the three models was different.

The third chemical was complex also for the QSAR models. CAESAR predicted the chemical as mutagenic with a high ADI value, SARpy predicted it as non-mutagenic with a low ADI and Toxtree predicted it as mutagenic with a low ADI. There were conflicting results and care should be applied in the overall evaluation, especially due to the low ADI for two models. CAESAR appeared to be more reliable. VEGA allowed the visualization of the similar compounds, which should always be evaluated: all the most similar chemicals are mutagenic and contain the 1,2-benzothiazole structure as the target compound does. However, all these similar mutagenic chemicals also contain an amino group on the benzene ring, which is absent in the target compound but is probably responsible for the mutagenicity. Thus, the evaluation of toxicity given by the model is only apparently supported by the similar compounds, because they seem to be mutagenic for the presence of an aniline moiety absent in the target compound. Thus, no solid conclusion can be obtained from the QSAR results. For this third chemical the evaluation based on the ToxRead seems more reasonable, probably because ToxRead is based on a larger database.

5. Conclusions

There is a need for more tools to improve predictions based on read across. Read across is gaining popularity since the introduction of more stringent regulations, such as the European REACH regulation.

Read across has been used more intensively than QSAR for registration, as reported by the European Chemicals Agency (ECHA) [35]. However, read across is typically performed by individual experts according to their own experience, usually on highly confidential data. No exercise has been conducted to check the reproducibility and the accuracy of this kind of read across, while for QSAR, some evaluations comparing the different programs have been made [24,25].

Read across requires experts in toxicology, chemistry, biology, environmental sciences and other fields. However, this kind of expert reasoning is rare, expensive and may also be subjective. Furthermore, experts may use different sets of rules and they may over-rely on past experience and miss new evidence. That is why expert reasoning may be irreproducible.

ToxRead aims to improve this current situation by providing a clear and objective basis which can be exploited by human experts in their analysis of the chemicals to be used for read across. The tool offers two basic resources: a set of rules/structural alerts which may explain the toxicity; and a similarity tool associated with a large database of chemicals with their properties. In this way the user is assisted in the navigation for the identification of the similar compounds through a set of pathways which represent the known toxicity processes.

Of course this kind of tool has to be flexible and should accept new sets of rules; currently we are increasing them for mutagenicity and we are also extending the tool towards other endpoints such as the bioconcentration factor and fish acute toxicity.

An open issue, as mentioned in the Introduction, is how to deal with uncertainty while making a read across prediction. Read across is not a probabilistic method and as such we cannot mathematically assign a measure of uncertainty. Nevertheless, some elements of uncertainty are indicated by ToxRead, as the accuracy of the structural alert (number of toxic chemicals) and also the similarity of the related chemicals, which is measured in a quantitative way. We have shown in our case studies how different structural alerts with different accuracy can be present in the same target molecule. Moreover, we have shown that exclusion rules are often stronger than toxicity rules. Combining all this together it is possible to attach to the result taken by the expert both the accuracy of the structural alert found and the similarity value of the most similar molecule considered to obtain numerical values expressing a sort of uncertainty.

For the similar chemicals, the highest similarity should drive the choice. Although, if the similarity of the chemical is lower than 0.75 its contribution may be questionable. Regarding the structural alert, the alert with a higher accuracy should be considered more useful; if the accuracy of the structural alert is close to 0.5, its relevance is modest. ToxRead is freely available and can be downloaded from its website [36].

Acknowledgements

We acknowledge the financial contribution of the LIFE programme, projects CALEIDOS and PROSIL. We are grateful to the NVIDIA academic partnership for providing the GPU nVidia Tesla™ C1060 board.

References

[1] E. Benfenati, G. Gini, S. Hoffmann, and R. Luttik, *Comparing in vivo, in vitro, in silico methods and integrated strategies for chemical assessment: Problems and prospects*, ATLA 38 (2010), pp. 153–166.

[2] E. Benfenati and R. Gonella Diaza, A. Cassano, S. Pardoe, G. Gini, C. Mays, R. Knauf, and L. Benighaus, *The acceptance of in silico models for REACH: Requirements, barriers, and perspectives*, Chem. Cent. J. 5 (2011), p. 58.

[3] D.J. Livingstone, *The characterization of chemical structures using molecular properties: A survey*, J. Chem. Inform. Comput. Sci. 40 (2000), pp. 195–209.
[4] OECD, *The OECD QSAR toolbox, version 3.1*, 2014; software available at http://www.oecd.org/chemicalsafety/risk-assessment/theoecdqsartoolbox.htm.
[5] *AMBIT, Version 2.4.11*, 2014; software available at http://ambit.sourceforge.net/.
[6] G. Patlewicz, D.W. Roberts, A. Aptula, K. Blackburn, and B. Hubesch, *Workshop Report: Use of 'read-across' for chemical safety assessment under REACH*, Regul. Toxicol. Pharmacol. 65 (2013), pp. 226–228.
[7] R. Benigni, C. Bossa, N.G. Jeliazkova, T.I. Netzeva, and A.P. Worth, *The Benigni/Bossa rulebase for mutagenicity and carcinogenicity - A module of Toxtree*, EUR-Scientific and Technical Report Series, EUR 23241 EN, Office for the Official Publications of the European Communities, Luxembourg, 2008.
[8] R. Benigni and C. Bossa, *Structural alerts for carcinogenicity, and the* Salmonella *assay system: A novel insight through the chemical relational databases technology*, Mutat. Res.-Rev. Mutat. 659 (2008), pp. 248–261.
[9] J. Kazius, R. Mcguire, and R. Bursi, *Derivation and validation of toxicophores for mutagenicity prediction*, J. Med. Chem. 48 (2005), pp. 312–320.
[10] H.S. Rosenkranz, Y.P. Zhang, and G. Klopman, *Studies on the potential for genotoxic carcinogenicity of fragrances and other chemicals*, Food Chem. Toxicol. 36 (1998), pp. 687–696.
[11] R. Benigni, T.I. Netzeva, E. Benfenati, C. Bossa, R. Franke, C. Helma, E. Hulzebos, C. Marchant, A. Richard, Y.T. Woo, and C. Yang, *The expanding role of predictive toxicology: An update on the (Q)SAR models for mutagens and carcinogens*, J. Environ. Sci. Health C 25 (2007), pp. 53–97.
[12] E. Benfenati, R. Benigni, D.M. Demarini, C. Helma, D. Kirkland, T.M. Martin, P. Mazzatorta, G. Ouedraogo-Arras, A.M. Richard, B. Schilter, W.G. Schoonen, R.D. Snyder, and C. Yang, *Predictive models of carcinogenicity and mutagenicity: Frameworks, state-of-the-art and perspectives*, J. Environ. Sci. Health C 27 (2009), pp. 57–90.
[13] B.N. Ames, *The detection of environmental mutagens and potential*, Cancer 53 (1984), pp. 2030–2040.
[14] J. Ashby, *Fundamental structural alerts to potential carcinogenicity or noncarcinogenicity*, Environ. Mutagen. 7 (1985), pp. 919–921.
[15] W.W. Piegorsch and E. Zeiger, *Measuring intra-assay agreement for the Ames* Salmonella *assay*, in *Statistical Methods in Toxicology, Lecture Notes in Medical Informatics*, L. Hotorn, ed., Springer-Verlag, Berlin, 1991, pp. 35–41.
[16] C. Hansch, P.P. Malony, T. Fujita, and R.M. Muir, *Correlation of biological activity of phenoxyacetic acids with Hammett substituent constants with partition coefficients*, Nature 194 (1962), pp. 178–180.
[17] J.A. Miller and E.C. Miller, *Searches for ultimate chemical carcinogens and their reactions with cellular macromolecules*, Cancer 47 (1981), pp. 2327–2345.
[18] J. Ashby and R.W. Tennant, *Chemical structure,* Salmonella *mutagenicity and extent of carcinogenicity as indicators of genotoxic carcinogenesis among 222 chemicals tested by the U.S. NCI/NTP*, Mutat. Res. 204 (1988), pp. 17–115.
[19] Q. Liao, J. Yao, and S. Yuan, *Prediction of mutagenic toxicity by combination of recursive partitioning and support vector machines*, Mol. Divers. 11 (2007), pp. 59–72.
[20] M. Zheng, Z. Liu, C. Xue, W. Zhu, K. Chen, X. Luo, and H. Jiang, *Mutagenic probability estimation of chemical compounds by a novel molecular electrophilicity vector and support vector machine*, Bioinformatics 22 (2006), pp. 2099–2106.
[21] T. Ferrari and G. Gini, *An open source multistep model to predict mutagenicity from statistical analysis and relevant structural alerts*, Chem. Cent. J. 4(Suppl. 1) (2010), p. 4.
[22] A. Perrotta, D. Malacarne, M. Taningher, R. Pesenti, M. Paolucci, and S. Parodi, *A computerized connectivity approach for analyzing the structural basis of mutagenicity in* Salmonella *and its relationship with rodent carcinogenicity*, Mol. Mutagen. 28 (1996), pp. 31–50.

[23] R.D. Snyder, G.S. Pearl, G. Mandakes, W.N. Choy, F. Goodsaid, and I.Y. Rosenblum, *Assessment of the sensitivity of the computational programs DEREK, TOPKAT and MCASE in the prediction of the genotoxicity of pharmaceutical molecules*, Environ. Mol. Mutagen. 43 (2004), pp. 143–158.
[24] N. Golbamaki Bakhtyari, G. Raitano, E. Benfenati, T.M. Martin, and D. Young, *Comparison of in silico models for prediction of mutagenicity*, J. Environ. Sci. Health C 31 (2013), pp. 45–66.
[25] A. Cassano, G. Raitano, E. Mombelli, A. Fernandez, J. Cester, A. Roncaglioni, and E. Benfenati, *Evaluation of QSAR Models for the prediction of Ames genotoxicity: A retrospective exercise on the chemical substances registered under the EU REACH regulation*, J. Environ. Sci. Health C 32 (2014), pp. 273–298.
[26] E. Benfenati, A Manganaro, and G. Gini, *VEGA-QSAR: AI inside a platform for predictive toxicology*, Proceedings of Workshop PAI 2013, 5 December 2013, Turin, Italy, pp. 21–28.
[27] E. Benfenati, G. Raitano, and G. Gini, *Overtaking barriers between toxicologists and computers: The example of VEGA*, Proceedings of IHCI 2013 –Multi Conference on Computer Science and Information Systems, Prague, 22–24 July 2013, pp. 268–272.
[28] G. Gini, T. Ferrari, D. Cattaneo, N.G. Bakhtyari, A. Manganaro, and E. Benfenati, *Automatic knowledge extraction from chemical structures: The case of mutagenicity prediction*, SAR QSAR Environ. Res. 24 (2013), pp. 365–383.
[29] The ANTARES project: Alternative Non-Testing methods Assessed for REACH Substances. Funded by the LIFE+ EC Programme 2010–2012. Website available at: http://www.antares-life.eu/.
[30] K. Hansen, S. Mika, T. Schroeter, A. Sutter, A. ter Laak, T. Steger-Hartmann, N. Heinrich, and K.R. Müller, *Benchmark data set for in silico prediction of Ames mutagenicity*, J. Chem. Inf. Model. 49 (2009), pp. 2077–2081.
[31] *Chemistry Development Kit*; software available at: http://sourceforge.net/projects/cdk/.
[32] *Java Universal Network/Graph Framework*; software available at: http://jung.sourceforge.net/.
[33] *HyperSQL*; software available at: http://hsqldb.org/.
[34] M. Floris, A. Manganaro, O. Nicolotti, R. Medda, G.F. Mangiatordi, and E. Benfenati, *A generalizable definition of chemical similarity for read-across*, J. Cheminf. in press.
[35] European Chemicals Agency (ECHA), *The Use of Alternatives to Testing on Animals for the REACH Regulation. Second report under Article 117(3) of the REACH Regulation*, ECHA, Helsinki, 2014.
[36] *ToxRead*, software available at: http://www.toxgate.eu/.