Computational Toxicology 29 (2024) 100300

Contents lists available at ScienceDirect

Computational Toxicology

journal homepage: www.sciencedirect.com/journal/computational-toxicology

Full Length Article

Developing and validating read-across workflows that enable decision making for toxicity and potency: Case studies with N-nitrosamines

Steven Kane, Dan Newman, David J. Ponting, Edward Rosser, Robert Thomas, Jonathan D. Vessey *, Samuel J. Webb, William H.J. Wood

Lhasa Limited, Granary Wharf House, 2 Canal Wharf, Leeds LS11 5PS, UK

ARTICLE INFO

Edited by Dr. Daniel Dietrich

Keywords:
Read-across
Nitrosamines
Software application

ABSTRACT

To reach conclusions during chemical safety assessments, risk assessors need to ensure sufficient information is present to satisfy the decision criteria. This often requires data to be generated and, in some cases, insufficient knowledge is present, or it is not feasible to generate new data through experiments. Read-across is a powerful technique to fill such data gaps, however the expert-driven process can be time intensive and subjective in nature resulting in variation of approach. To overcome these barriers a prototype software application has been developed by Lhasa Limited to support decision making about the toxicity and potency of chemicals using a read-across approach. The application supports a workflow which allows the user to gather data and knowledge about a chemical of interest and possible read-across candidates. Relevant information is then presented that enables the user to decide if read-across can be performed and, if so, which analogue or category can be considered the most appropriate. Data and knowledge about the toxicity of a compound and potential analogues include assay and metabolism data, toxicophore identification and its local similarity, physico-chemical and pharmacokinetic properties and observed and predicted metabolic profile. The utility of the approach is demonstrated with case studies using *N*-nitrosamine compounds, where the conclusions from using the workflow supported by the software are concordant with the evidence base. The components of the workflow have been further validated by demonstrating that conclusions are significantly better than would be expect from the distribution of data in test sets. The approach taken demonstrates how software implementing intuitive workflows that guide experts during read-across can support decisions and how validation of the methods can increase confidence in the overall approach.

1. Introduction

Read-across methods can be used for the prediction of many chemical properties. The basic concept is to gain insight into one or more properties of a target compound by comparison with the known properties of other compounds, usually termed analogues. The method is particularly useful for making a prediction of toxicity where there are insufficient data for a compound and where testing may be expensive, time consuming or be unacceptable for ethical or legal reasons.

There is currently no standard method for performing a read-across analysis for toxicity though several approaches have been published over the last few years [1–9]. Consistency of read-across methodology, or workflow, is extremely important in assessing a read-across analysis and regulatory agencies, who may be presented with read-across analyses among submission data that they have to consider, have commented on this [10–12] with the European Chemicals Agency (ECHA) publishing their own Read-across Assessment Framework (RAAF) [1]. Read-across has also been used alongside other New Approach Methodologies (NAMs) [13–15]. All approaches to read-across have common elements based on deciding which compounds form close analogues to

Abbreviations: AI, acceptable intake; AOP, adverse outcome pathway; API, active pharmaceutical ingredient; CPCA, Carcinogenic Potency Categorisation Approach; DNA, Deoxyribose nucleic acid; ECHA, European Chemicals Agency; EMA, European Medicines Agency; LCDB, Lhasa Carcinogenicity Database; MRE, mean relative error; NAM, new approach methodology; NDEA, nitroso diethyl amine; NDMA, nitroso dimethyl amine; OECD, Organisation for Economic Co-operation and Development; RAAF, Read across assessment framework; SMILES, simplified molecular input line entry system; TD50, dose at which 50% of subjects show tumor formation; US EPA, United States Environmental Protection Agency.

* Corresponding author.
E-mail address: jonathan.vessey@lhasalimited.org (J.D. Vessey).

https://doi.org/10.1016/j.comtox.2024.100300
Received 20 November 2023; Received in revised form 12 January 2024; Accepted 24 January 2024
Available online 29 January 2024

the target, based on their chemical structure [5] and other factors such as physico-chemical properties or metabolism [16–18] which are considered to be important to the toxicity of interest. Deciding on what constitutes 'similar' in terms of chemical structure is not straightforward and measures of similarity which focus on a specific area of the target molecule may be more successful than measures of similarity of the whole target molecule [5,19,20].

A read-across analysis may produce a binary toxic/non-toxic outcome or a potency expressed as a point value or range [21]. Read-across may reference a single compound or, alternatively, a category approach may be taken [14,22–24], in which a well-described group of compounds with structural [22,25] and possibly physico-chemical properties [26] is defined and the toxicity of the target is considered in relation to this group.

Read-across requires a number of individual decisions each of which may materially influence the final conclusion and therefore a software solution that systematically identifies and simplifies these decisions offers significant benefit. To this end, several software solutions have been made publicly available. The United States Environmental Protection Agency (US EPA) have implemented a workflow called GenRA, generalised read-across, in the CompTox chemical dashboard [27]. GenRA attempts to make the read-across process algorithmic, and thus reproducible. It does this by using structural descriptors (Morgan fingerprints and torsion descriptors) as well as bioactivity descriptors. Benfenati and co-workers have reported ToxRead [28] which is aimed specifically at read-across for mutagenicity. It uses structural alerts and other relevant features to identify analogues of a target chemical. The OECD toolbox supports a read-across workflow based on category and sub-category definitions [11]. Kutsarova *et al.* have published an automated method of using the QSAR toolbox for prediction of acute oral toxicity [29,30]. A simple workflow for read-across implemented in Knime has been demonstrated for aromatase activity [31] and a framework which allows for different toxicity endpoints has been published by Moustakas *et al.* [32] An approach to quantifying the suitability of analogues for toxicity assessment has been published [33] the approach goes beyond using simple structural similarity for selection of analogues: biological, physico-chemical and metabolic properties are taken into account when generating an overall score for potential analogues. Commercial tools have also been developed: QSAR Flex is a tool developed by MultiCASE which also has a module to aid in read-across [34], and Instem have recently added read-across support to their Model Applier Computational Toxicology software, though few details are available at the time of writing [35]. Commercial tools may have the advantages of connecting to proprietary data held by the users or to data sharing initiatives which may themselves have restricted access.

In this paper we report how we have integrated components into a prototype application which supports the user to assemble all the relevant data, interact with different facts during read-across and make an overall read-across assessment. We illustrate the decision making process with N-nitrosamine examples and consider more broadly the how *in silico* frameworks can support more general read-across approaches. The emphasis of the prototype application is to support the human expert, suggesting a standard mode of use, but allowing the expert to apply their expert knowledge.

Nitrosamines and their toxicity have come to the fore in recent years as their presence as impurities in widely used pharmaceuticals [36] has led to withdrawal of drugs impacting on the healthcare of people reliant on the drugs as well as the cost to pharmaceutical companies. Initially, nitrosodimethylamine (NDMA) or nitrosodiethylamine (NDEA) were discovered as impurities as a result of manufacturing processes, but nitrosated derivatives of the API have also triggered recalls[37].

The most concerning aspect of the toxicity of nitrosamines is their carcinogenic potential, particularly as some of the drugs concerned are expected to be taken over a lifetime. As some nitrosamines are known to be very potent carcinogens, there is a need to assess the risk associated with these compounds without experiments on animals. By using existing data for compounds of this class, it might be possible to agree an acceptable intake (AI) for novel, untested, nitrosamines.

The reader should note that classification of chemicals as carcinogen/non-carcinogen has been questioned [38], indeed the term 'cancer hazard' rather than 'carcinogen' has been recommended. It is outside the scope of this paper to discuss terminology, and we use the term 'carcinogen' only since it is widely used in the scientific community at present. Limits of exposure to some nitrosamines have been derived based on read-across from NDMA or NDEA [39,40] which are both known to be extremely toxic compounds and thus have AIs which are very low, 96 ng/person/day for NDMA and 26.5 ng/person/day for NDEA, although these values have been questioned [41]. Many other nitrosamines are known to be carcinogenic in mammals, though their toxicity is generally lower than NDMA while still other nitrosamines do not exhibit carcinogenicity [42]. A read-across approach to estimating daily limits based on nitroso piperidines has recently been published [43]. The European Medicines Agency has recently augmented its advice on nitrosamine impurities [44] with a Q & A document refining its approach to AI values of nitrosamines [45] The EMA guidance includes a decision tree for AI but other techniques, including analogues selected by read-across, are considered.

We illustrate how the application can be used to identify toxic nitrosamines and assign a potency value to them expressed as the TD_{50} in mg/kg/day (dose causing tumour in 50 % of the exposed individuals). The approach uses both defined categories and focussed similarity as well as taking into account metabolism (either observed or predicted) and calculated physico-chemical properties and biological clearance. The components within the software established to support these expert-driven assessments were also validated to support further justification of the approach.

2. Methods

2.1. Data

Toxicity data from several databases were used. Data from Kaptis (in-house version) [46] were augmented with data from the Vitic database (version 2022.1.0) [42,46,47] and freely available Lhasa Carcinogenicity Database (LCDB) version 2021.1.0. [46,47]. The Kaptis data consisted of overall calls for relevant assays and some more detailed data about the observations made as part of each assay where they were available. Vitic and LCDB contributed observations for specific assays and derived potencies as TD_{50}.

Although all the data in the constituent databases were available for comparison by the prototype application, in practice, only compounds with the *N*-nitroso group were considered as possible analogues. In total there were 508 *N*-nitrosamines with some data relevant to forming a conclusion about either carcinogenic activity or potency.

The set of 508 compounds was compiled separately for the validation study. The breakdown of this dataset was as follows. The data from Kaptis yielded 244 compounds with either an Ames call (202 positive, 33 negative, 9 other) or a rodent carcinogenicity call (68 positive, 9 negative, 2 other). Vitic yielded 434 compounds with an overall call for Ames (346 positive, 78 negative, 10 other) and 197 compounds with a call for an assay in the carcinogenicity table. LCDB gave 140 compounds with a call of which 48 had a TD_{50} (either Gold TD_{50} or Lhasa TD_{50} or both) [48]. The different sources were combined on a conservative basis, i.e. any positive call resulted in the compounds being considered positive and compounds were considered negative if the only data available were negative calls. Compounds were removed if they had equivocal or conflicted calls but no positive calls. This left a set of 494 *N*-nitrosamines: 418 positive and 76 negative. A subset of these data was compiled where compounds which had only Ames data were excluded. This gave a set of 250 compounds with a carcinogenicity call: 201 positive and 49 negative.

The Lhasa Limited metabolism database provided data about

metabolites for compounds where these were known [49].

2.2. Knowledge

The links between assays and adverse outcome in the context of expert derived AOPs were provided by Kaptis [46,50].

Forty-six structural features considered to be of potential significance in the classifying of toxicity or evaluating carcinogenic potency of N-nitrosamines have been evaluated in-house and collaboratively [25,51–53]. They are listed in more detail in Thomas *et al.* [51], which had 41 patterns, and four patterns have been added for completeness of coverage of chemical space; these do not substantively affect the results determined by Thomas *et al.* Further refinements in light of the EMA Q & A document were made by Ponting [53]. These features can be used to develop potency categorisation where features lead to statistically significant increases, or decreases, in potency [51,53].

Toxicophore identification was provided by a modified web service of Derek Nexus [46,47,54] using the 2022.1.0 version of the knowledge base.

2.3. Predictions

Metabolite and site of metabolism predictions were provided by an in-house web service using transformations derived from Meteor Nexus version 3.1.0 [54,55].

An in-house model of clearance was used. The dataset of human clearance values used for model training contained 1325 compounds from the ChEMBL database version 25 [56]. The dataset was standardised such that clearance units were ml min^{-1} kg^{-1}. A random forest was built using Python (3.8.8), Scikit-learn (1.0.2) [57] and RDKit (2021.09.2) [58]. Descriptors based on physico-chemical properties were calculated from compound SMILES strings using RDKit's "Molecular Descriptor Calculator". The first 123 descriptors were used for model building. Model hyperparameters were optimised using a 5-fold cross validation approach with negative mean squared logarithm base 10 scoring. The final random forest model contained 1000 estimators with a maximum depth of 50 and a maximum number of features set at 40. 20 % of the original dataset was withheld and used as a test set in order to assess the final performance of the model. The average model performance based on the test set was 0.495 root mean squared log (base e) error with a root mean squared error of 12 ml min^{-1} kg^{-1}. 95 % confidence intervals were calculated using conformal prediction based on the normalised nonconformity measure (equation 32) from Papadopoulos *et al.* [59]. This also required the use of a nearest neighbour model (scikit-learn 1.0.2) which was built from a dataset withheld from training the underlying random forest clearance model. The nearest neighbour model was based on 5 nearest neighbours via a ball tree algorithm. The applicability domain was defined as those compounds within the range of minimum and maximum values of each descriptor as observed in the training set.

An in-house model of solubility was used [60].

Log P values were provided by JPLogP [61].

2.4. Software

The prototype application was developed using an internal suite of tools [50,62]. The data, knowledge and predictions in the preceding sections were provided as web services that were called by the prototype application. To enable batch processing for validations the 45 categories in Cross and Ponting [25] and the focussed similarity method were incorporated into a Knime workflow.

3. Results

3.1. Workflow supported by the software

The workflow supported by the prototype application is based on the ECHA RAAF and is shown in Fig. 1; here it is described in more detail. The reader is also pointed to the examples which illustrate the approach in a less abstract form.

3.1.1. Step 1: Submit query compound (target); choose endpoint

The compound to read-across to – the query compound or target – is entered into the prototype application, either drawn by the user or

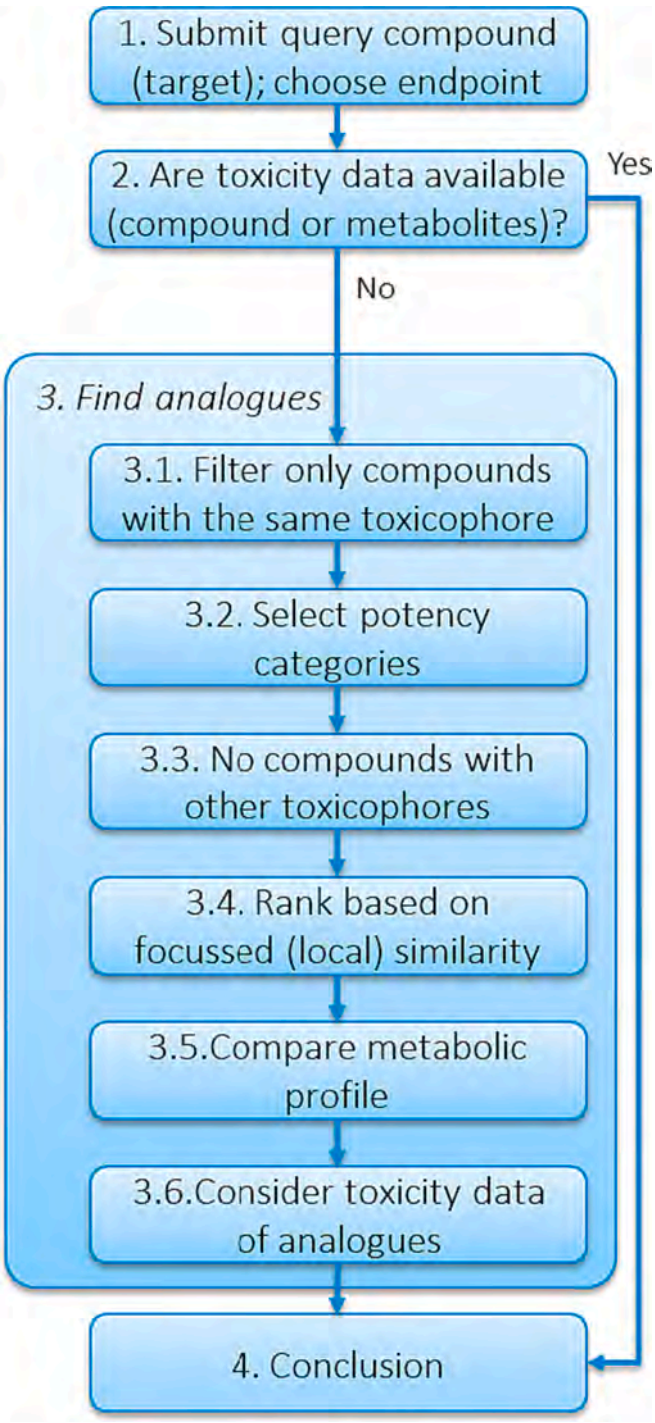

Fig. 1. Workflow supported by the prototype application.

imported as a SMILES or MOLfile. Standardisation of the structure information is essential to ensure consistent representation and relevant data is retrieved from databases and prediction systems. In the prototype, the method(s) of standardisation can be selected by the user, but the recommended default is for the structure to undergo validation (looking for valency violations etc.) and normalisation (transforming to a favoured tautomer or mesomer). Contextualisation (removing salts and dealing with mixtures; removing stereochemistry etc.) might be appropriate, but, when considering similarity beyond simple chemical structure, e.g. solubility or other physicochemical effects, this may not be suitable. In-house tools using a rule-based approach are used for these steps [50]. If the structure is found in the databases, then the prototype application will display this to the user.

To frame the safety assessment the user also selects an endpoint or adverse outcome for the read-across. Whilst this is optional, so allowing the user to do a very general read-across, it is recommended to select a toxicity endpoint to make most use of the methods of narrowing down potential read-across analogues. In this instance the underlying toxicity knowledge is based on adverse outcome pathways (AOPs) and the Kaptis knowledge base, which links toxicology assays to AOP key events and adverse outcomes [63,64]. The user can select any subset of the AOPs which lead to an adverse outcome if known pathways are the focus of the read-across. The linking of assay data, from which the read-across might be made, to toxicological endpoints/adverse outcomes using AOPs allows the user considerable control in selecting appropriate assays and findings. For example, the relevance of animal data to human cancer hazard is disputed [65,66] and certainly in cases where animals are found to develop tumours which have no equivalent in humans [67] the user may consider the animal data to have no relevance to a human safety assessment.

3.1.2. Step 2: Are toxicity data available?

In keeping with the ECHA RAAF, any available toxicity data about the submitted compound are retrieved from the databases. Firstly, data about the submitted compound are presented; any available data are organised by the assay and the measurement – which are selectable by the user – and the observations are presented in a table form and appropriate graphical format e.g. histogram.

This may be sufficient for the user to assess the toxicity of the submitted compound. However, further relevant data may be available. In particular, and again as suggested by the RAAF, any known metabolites of the submitted structure are retrieved from a database of published metabolism studies. Available toxicity data for each of the identified metabolites are also retrieved from the toxicity databases and displayed in the same way as for the submitted compound. Knowing the metabolic fate of the submitted compound may reveal that it is converted into a compound for which sufficient data are available to allow the user to make a decision on the toxicity of the submitted compound.

For novel structures, it is unlikely that toxicity or metabolism data will be available in the public domain, though it may well be that users with private or proprietary data will have relevant information.

3.1.3. Step 3: Find analogues

Assuming that either no data are available for the submitted compound directly or that such data that are available are insufficient to come to a conclusion about its toxicity, the next steps involve the read-across process in order to identify suitable analogues. This part of the workflow incorporates the steps in the RAAF for considering both a category approach and finding individual analogues.

This section is divided into the individual steps shown in Fig. 1.

3.1.3.1. Step 3.1: Filter only compounds with the same toxicophore.

The first step in the read-across process is to identify the parts of the target compound that are relevant to the toxicity endpoint selected. The user might already have prior knowledge of the toxicophore that they are most interested in, but a more robust approach is to use structural alerts from knowledge base systems. In this instance toxicity predictions from Derek Nexus are used to identify all the relevant toxicophores in the target molecule.

In terms of making a read-across assessment it is important that candidate read-across analogues have a similar mechanism of action to the target compound. Therefore, only considering compounds with the same toxicophore(s) as the target is important. A candidate analogue which did not contain the identified toxicophore would be a poor choice, because any toxicity it exhibited would likely arise from a different mechanism, although even then there may be some subjectivity in this decision depending on how the toxicophores are implemented.

Where the toxicophore is unknown, it is obviously not possible to make a read-across of positive toxicity that is as convincing as where the toxicophore is known. Furthermore, in order to read-across a negative toxicity, i.e. to read-across from a compound considered to be similar which also does not contain a toxicophore and is found to be non-toxic, it is important that the user is presented with similar compounds which nevertheless do contain a toxicophore, if only to be able to discount them as potential analogues. The prototype application supports this use case by allowing the user to include in the assessment compounds which contain toxicophores not present in the target.

3.1.3.2. Step 3.2: Select potency categories.

The use of potency categories has been identified as a valuable step in the read-across process. In some cases, a category can be quite tightly defined and is exclusive of other chemicals, in other cases potency categories might overlap.

In the case of nitrosamines, potency categories have been proposed by Cross and Ponting [25] and refined in further publications [51–53]. These categories share many features with the Carcinogenic Potency Categorisation Approach (CPCA) published by the EMA [68] which classifies *N*-nitrosamines into five different categories with four different AI values. Potency categories for nitrosamines have also been proposed by Dobo *et al.* [69]. The prototype application uses a set of structural categories which are related to potency categories and, typically, any particular nitrosamine will fall within the definition of several of these structural categories. To handle these subtleties around chemical similarity is complex and careful user interface design is required so that multiple perspectives of chemical similarity can be examined by the user. One way that the prototype application handles this is through interactive plots such as Venn diagrams and UpSet plots that can display the number of chemicals that contain each combination of structural categories. Diagrams such as these can inform which analogues share important features which can be considered for read-across. However each diagram can have their own strengths and limitations depending on the complexity of the scenario. For example, Venn Diagrams can get unwieldly with more than 5 categories and therefore the user will have to prioritise which features are important for consideration. The reader is referred to the examples below where this situation is described in more detail.

3.1.3.3. Step 3.3: Removal of compounds with other toxicophores.

When considering candidate analogues for read-across of a particular toxicity endpoint, it may be sensible to remove compounds which can cause the same toxicity via mechanisms not available to the target compound, that is to say remove compounds with other relevant toxicophores. This is because if there are other mechanisms leading to that endpoint and affecting the potency, the analogues could be considered less relevant. However, in the case where one toxicophore is known to be very potent and a second toxicophore is known to be less so, a read-across from an analogue with other toxicophores might be acceptable.

The prototype application gives the user the option to filter out compounds which fire Derek Nexus alerts for the endpoint of interest which are not fired by the target compound. The user can refine this to allow compounds which do indeed fire other relevant alerts and this

may be useful where the toxicophore is defined by different alerts that overlap.

3.1.3.4. Step 3.4 rank based on focussed (local) similarity. After the user has narrowed down the possible read-across candidates by structural category and relevant toxicophore(s), the remaining compounds can be ranked by similarity to the target. Many chemical similarity measures use a whole molecule similarity but looking at global similarity alone is not appropriate in the case of selecting an analogue for read-across purposes because the local environment around the toxicophore plays a strong role in affecting the activity of the target.

The prototype application uses a focussed similarity approach, whereby the atoms of the toxicophore (or other feature of interest) are selected as a core and the user selects which toxicophore (if there are more than one, or more than one occurrence) and how far away – in terms of bonds – the substructure for the similarity measure is. The equivalent substructure is selected from each of the candidate analogues by identifying the equivalent core (for example the atoms that match the same toxicophore identified by a Derek Nexus alert). From the substructures a fingerprint is generated – in the case of the prototype application this is a Ceres fingerprint [62] – and the similarity of candidates is measured as the distance between the candidate analogue focussed fingerprint and that of the target compound. The measure of the distance is selectable by the user; typically, a Tanimoto distance between the fingerprints is used but the option of Tversky distances is also given. In practice, Tversky distances between the fingerprints are smaller, indicating greater similarity, where there are no other chemical functions in the substructures, thus indicating 'simpler' compounds as more similar to the target compound.

Given that the distance between the fingerprints is a function of the size of the substructure in terms of number of bonds' distance from the focus of similarity, the actual value of the similarity measure is less important than the ranking of the candidate analogues.

3.1.3.5. Step 3.5 Compare metabolic profiles. An important step in making a read-across assessment as described in the ECHA RAAF, is to consider the metabolism of the target compound and how it compares to a suggested analogue. Clearly, if a target compound and proposed analogue differ in the metabolic transformations expected or observed around the toxicophore, the candidate would be considered a less convincing analogue. Even where any expected or observed metabolism is remote from the toxicophore it may still have a significant impact on the solubility, distribution or excretion of the target or candidate analogue and may be a factor in augmenting or undermining the user's confidence in the appropriateness of a candidate as an analogue.

In the prototype application, the target compound and candidate analogues have the observed or predicted metabolisms compared. If there are data on the target compound or candidate analogues, then it can be retrieved from the metabolism database. In the more common situation for novel compounds where there are no metabolic data, the prototype application can call on a web service implementation of metabolism prediction based on the knowledge in Meteor Nexus. The metabolism predictions used by the prototype application are limited to single generation metabolic transformations in a general mammalian model. The limiting of the predictions to a single generation of transformations is for two reasons: firstly, and principally, because it emphasises the site in the target molecule and potential analogues at which metabolism takes place, and secondly, because it keeps the number of generated metabolites to a reasonably manageable number. The Meteor Nexus prediction model is a general mammalian one, though the user can change the prediction settings to utilise species specific predictions of metabolism where they are known. In contrast to metabolism studies where quantitative metabolic data are available, Meteor Nexus is only able to give a score [70] as an indication of likelihood of formation of a metabolite; the score is reported to the user and is used in ordering how the predicted metabolites are viewed.

The prototype application displays the predicted metabolites of the target compound and the metabolic biotransformations that are predicted to arise in a table form alongside similar predictions for the remaining candidate analogues. In addition, the user can compare the metabolism of the target compound with that of another arbitrary structure. This is particularly useful in the case of N-nitrosamines where the user might want to look at the predicted metabolism of the parent amine alongside that of the nitrosamine derivative, to see if there are biotransformations specific to the nitrosamine that would imply a different general profile to that of the parent amine.

3.1.3.6. Step 3.6 Consider toxicity data of analogues. Having reduced the number of candidate analogues using structural categories and focussed similarity, the user can retrieve the relevant toxicity data for them. In the prototype application, this is presented both in aggregated form using graphics such as histogram and tabular form allowing the user to investigate both individual candidate analogues and a set of analogues at the same time.

3.1.4. Conclusion

The different pieces of information in the read-across workflow are summarised in the prototype application in a table form. The most similar candidate analogues are ranked by their focussed similarity and presented with any relevant toxicity data. In addition, other calculated parameters which may contribute to the overall assessment are presented; these include molecular weight, solubility, log P and calculated clearance. These are presented graphically in a stacked bar chart, coxcomb plot, skyline plot, radar plot or heatmap with the physicochemical and other properties normalised on the difference between the value for the candidate analogue and the target compound. The candidate analogues may also be ranked on the total over all the similarity measures (focussed similarity and other properties). The inclusion of other properties ensures that global similarity is also taken into account when making the read-across assessment. Having the ability to select the most relevant analogues, rank them appropriately and displaying properties in a user-friendly manner provides the foundations for experts to confidently perform read-across assessments.

3.2. Examples

The workflow is illustrated by the following examples.

3.2.1. Example 1: N-nitroso varenicline

N-Nitroso varenicline is an impurity found in the smoking cessation medication varenicline. The presence of *N*-nitroso varenicline at significant levels above the permitted limit led to some recalls of batches of the drug in 2021 [52,71].

The structure of *N*-nitroso varenicline is entered into the prototype application along with that of varenicline as a secondary structure. Malignant neoplasm (i.e. carcinogenicity) is selected as the adverse outcome/endpoint and the relevant AOPs in the database are presented to the user. The prototype thus highlights a difference in the terminology between the AOPs (malignant neoplasm) with respect to that of the toxicology endpoint (carcinogenicity); through the use of controlled vocabulary, user-friendly tools can provide consistent representation in the interface. As *N*-nitroso varenicline is not in the underlying database no standardised versions of the structure of the compound are shown, see Fig. 2.

As there are no relevant toxicity or metabolism data for *N*-nitroso varenicline the user moves on to the read-across process. The target compound fires Derek Nexus alert 070 for carcinogenicity of *N*-nitro and *N*-nitroso compounds and the prototype can query the underlying database to connect alert 070 to a key event in the AOP 'Electrophilic reaction with DNA leading to carcinogenicity'. There are 13,433

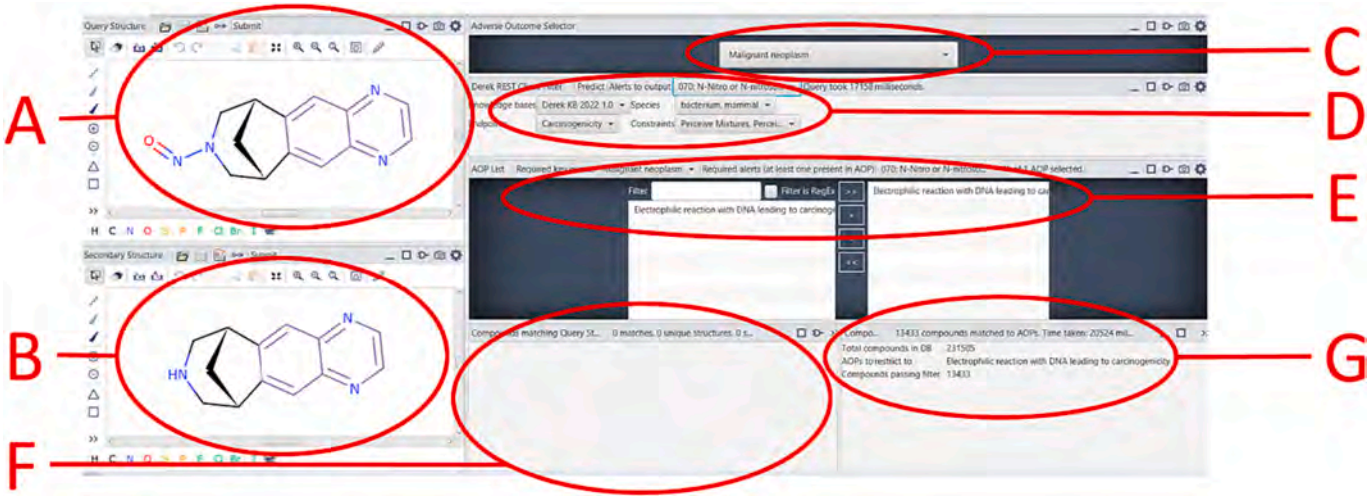

Fig. 2. The prototype application compound entry and endpoint selection screen. (A) The target compound N-nitroso varenicline. (B) Optional secondary structure varenicline. (C) Selected adverse outcome/endpoint. (D) Derek Nexus prediction parameters and identified toxicophores (E) corresponding AOPs in Kaptis. (F) Any standardised version(s) of the target structure known to the database(s). (G) Compounds in the selected database(s) for which there are data associated with assays reporting for key events in the selected AOP.

compounds in the underlying database with data relevant to the AOP which are retrieved. The search for analogues can then be limited to compounds which also fire alert 070, which returns 480 more pertinent examples.

The next stage is to consider the structural categories. *N*-nitroso varenicline falls into five of the defined structural categories in the prototype application, *viz.* has an alpha CH_2; generic nitrosamine; piperidine derivative and *N*-nitrosamine in rings of sizes six and seven. There are no compounds in all five of these categories and the user can select all five categories to generate a Venn diagram to investigate how the 480 candidate compounds can be reduced to more relevant ones as shown in Fig. 3, where the user has also filtered out compounds which fire other Derek Nexus carcinogenicity alerts. Of the different overlap regions of the Venn diagram, there are no compounds in the centre section and only one of the 4-out-of-5 overlaps contains structures. In the populated 4-out-of-5 segment there are twenty-six compounds; these do not contain a ring of size seven. The prototype application lets the user select any number of sectors and the associated compounds appear in a structure matrix for inspection.

The compounds in the matrix in Fig. 3 are unordered, but it can be seen that they have a high degree of structural homogeneity, and there may already be enough information to come to a conclusion on the toxicity of the target. The categories used by the prototype application are linked to potency values as proposed by Thomas *et al.* [51] and further refined by Ponting [53]. This uses a bucketing algorithm based on the potency categories listed in Table 1.

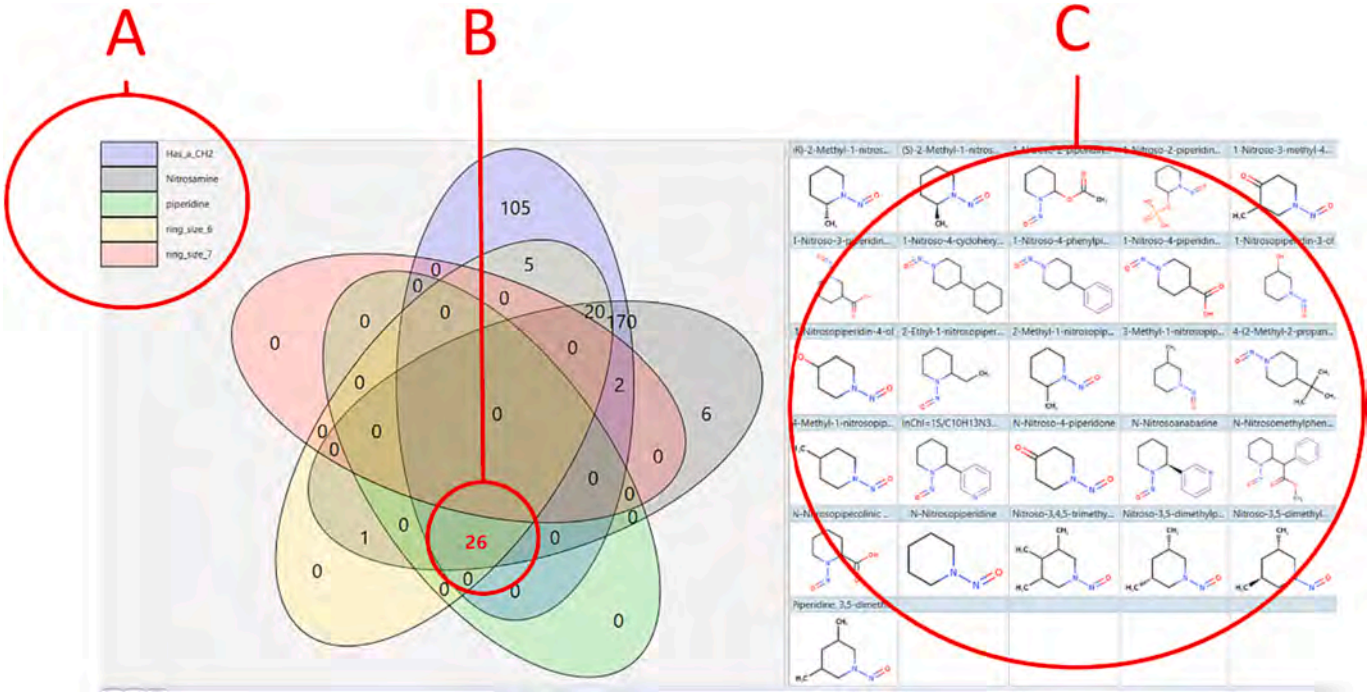

Fig. 3. Selected structure categories and selected overlap with the associated compounds in a Venn diagram in the prototype application. (A) Selected structure categories as a key to the Venn diagram. (B) Selected region of overlap. (C) Structures contained in the selected regions of overlap.

Table 1
Potency features after Thomas et al. [51] modified by Ponting [53].

Potency increasing features	Potency decreasing features
One or both α-carbons are ethyl or methyl	Carboxylic acid group anywhere
Benzylic, allylic or propargylic conjugation	One or both α-carbons are aromatic
Weak β-electron withdrawing group, e.g. carbonyl	Strong β-electron withdrawing group, e.g. -CF_3
One or both β-carbons are methyl	Steric hindrance, e.g. isopropyl group
Has two hydroxyl groups positioned β and γ	Lack of α-hydrogen, e.g. *tert*-butyl group
	Has one or more β-hydroxyl groups
	Both sides of the nitrosamine are in large chains (five atoms or more).

The bucketing algorithm that is proposed in Thomas *et al.* [51] states that if a compound contains features from neither list it is considered medium (150–1500 ug/kg/day); if it contains features from both lists it is also considered medium but requires further expert review; if it contains only features from the increasing list, then it is considered strong (18–150 μg/kg/day) and if it contains only features from the decreasing list then it is considered weak (>1500 μg/kg/day). On this basis, the target compound contains features from neither list so the user would conclude that the target is in the medium potency category i.e. should be limited to 150–1500 μg/kg/day.

Assuming that the user wishes to investigate further, they can move on to the next stage of the read-across process by returning to the twenty-six candidate analogues selected in Fig. 3 and ranking them using a focussed similarity to the target. The prototype application allows the user to select the toxicophore from which the focussed similarity is to be measured, as shown in Fig. 4.

The target compound in Fig. 4 is shown with explicit hydrogens as they may form part of the toxicophore. The user has selected the relevant Derek Nexus alert and the occurrence of the alert in the target compound is highlighted. In some molecules, there may be more than one alert and/or may be more than one match of an alert and the prototype application allows the user to select which of the identified toxicophores to consider as the focus of the similarity. As can be seen the user can select the difference measure between the fingerprints of the target and

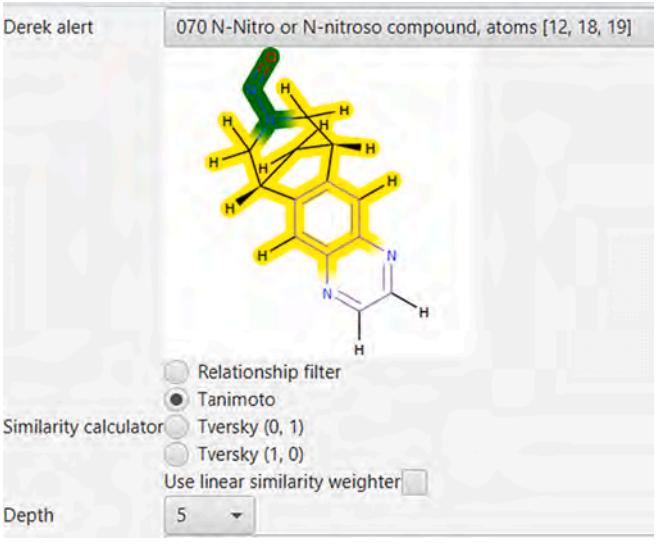

Fig. 4. Toxicophore match of the Derek alert 070 highlighted in green on the target compound and options for measuring the focussed similarity. The distance from the toxicophore that is included in the focussed similarity is selected with the Depth parameter and highlighted in yellow. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

candidate analogues (Tanimoto, Tversky (1,0) or Tversky (0,1)) as well as how many atoms and bonds away from the focus should be considered when generating the fingerprints to measure the focussed similarity (indicated by the Depth setting). In Fig. 4 the depth setting is 5 so the fingerprint will be generated using atoms at 5 bonds or less distance; this effectively excludes the remote heteroaromatic ring from consideration and this is also indicated in the Figure.

The twenty-six analogues are ranked according to the focussed similarity settings and displayed to the user; in Fig. 5 the three compounds with the highest focussed similarity are shown. The highlighting in Fig. 5 is calculated using an approach based on circular fingerprints [62,72], with areas of the target and analogue considered structurally similar shown in green and areas considered dissimilar shown in red. This highlighting is independent of the focussed similarity calculation but is nevertheless useful for the user in identifying regions of the molecules which are structurally different. Note that the visual similarity is based on a different approach to that of Rinker and Landrum [73]. The prototype application attempts to orient the analogues relative to the target using the focus of similarity as a reference point.

The three candidate analogues in Fig. 5 are, of course, isomers: two are stereo-defined and the third does not have defined chemistry so is assumed to be a mixture. A more sophisticated similarity measure would consider the compound with *cis* methyl substituents as slightly more similar to *N*-nitroso varenicline than the others as they match the positioning of the aromatic carbons in the target structure – though in this case there is no difference in the toxicity reported for the different isomers.

With the candidate analogues ranked, the user is directed to the comparison of the predicted metabolism. Parts of the predicted metabolisms for *N*-nitroso varenicline, the secondary structure of varenicline and the highest ranked candidate analogue are shown in Fig. 6. The prototype software orders the predicted biotransformations from left to right by their score [70] in the metabolism prediction web service for the target compound and then lexicographically, though the user can rearrange the prediction for ease of viewing. Fig. 6 shows that there are biotransformations unique to the secondary structure and that the highest ranked analogue and the target undergo equivalent metabolism of oxidation beta to the nitrosamine by two different biotransformations – thus augmenting the suggestion of the highest ranked candidate as a good analogue for the target compound.

Revisiting the comparison of target and the analogue, the prototype application allows the augmentation of the side-by-side similarity shown in Fig. 5 with the sites of metabolism calculated by the metabolism web service and corresponding to the biotransformations shown in Fig. 6; this is shown in Fig. 7.

The final step of the read-across analysis brings together the toxicity data for the candidate analogues as well as calculated physicochemical and biochemical properties of the candidates and the target. In any particular read-across, different biochemical and physiochemical properties will be of interest and the prototype application gives the user access to many different calculated properties.

In Fig. 8 five of the most similar candidate analogues are shown alongside the target compound together with the relevant data summarised with some calculated physico-chemical data and the prediction from an in-house clearance model. TD_{50} values are shown from the LCDB and are in mg kg^{-1} day^{-1}; the predictions from the clearance model are shown with 95 % confidence intervals. Where Ames data has been retrieved from Vitic the prototype application applies red-amber-green colour coding to the experimental data according to how it complies with regulatory guidance [74]: the guideline compliance is considered strong and coloured green if a positive result, even if in a single-strain, is obtained, or if the compound is found to be negative in the recommended combination of 5-strains +/− S9. If a compound is found to be negative in 4-strains +/− S9 it is considered to be acceptable in most cases, rated medium and coloured amber. Other cases are considered weak and coloured red. The combination of strains

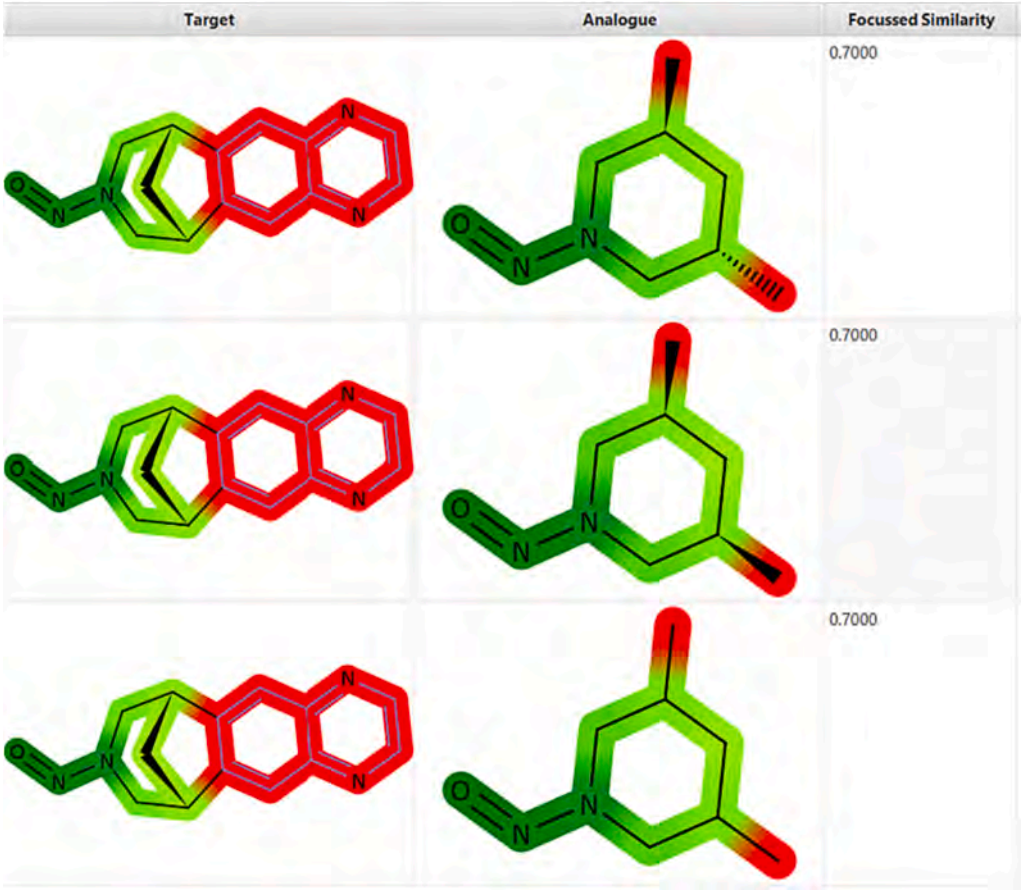

Fig. 5. Top-ranked candidate analogues for N-nitroso varenicline showing the structural similarity highlighting.

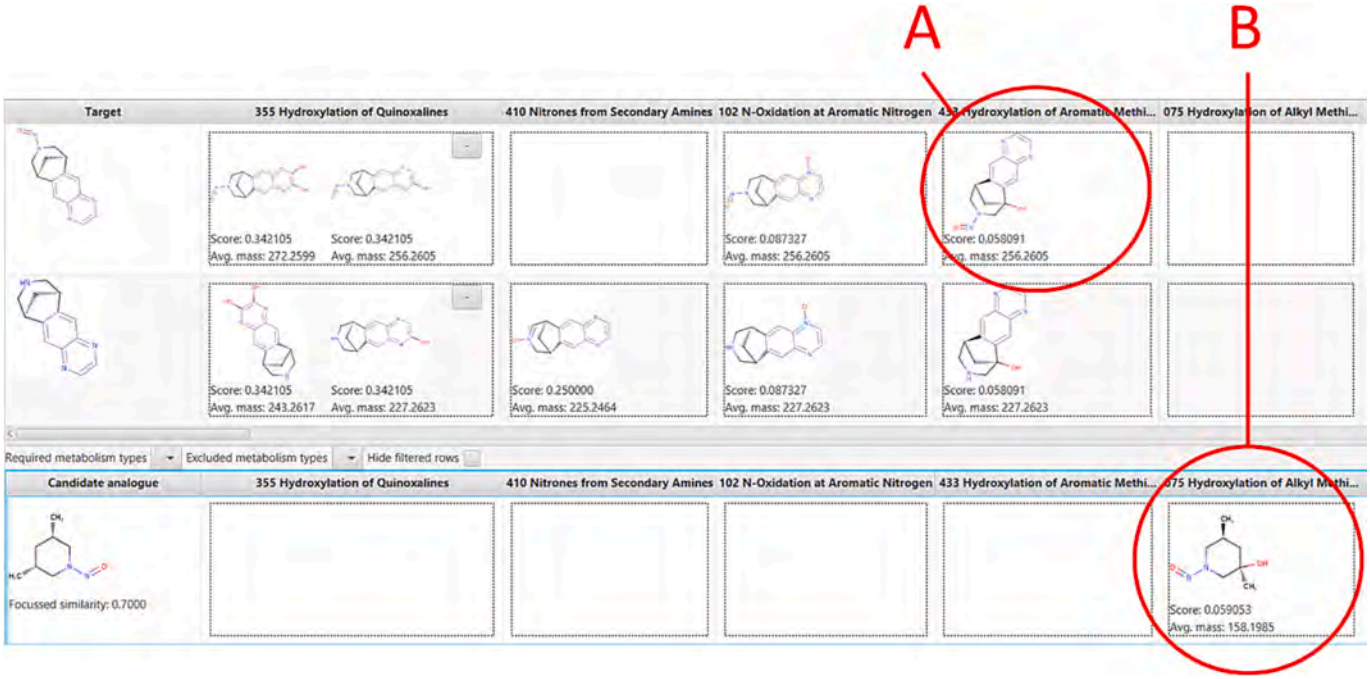

Fig. 6. Extract from the predicted metabolism of N-nitroso varenicline, varenicline and the highest ranked candidate analogue of N-nitroso varenicline. Biotransformations marked (A) and (B) show similar derivatives obtained by different biotransformations.

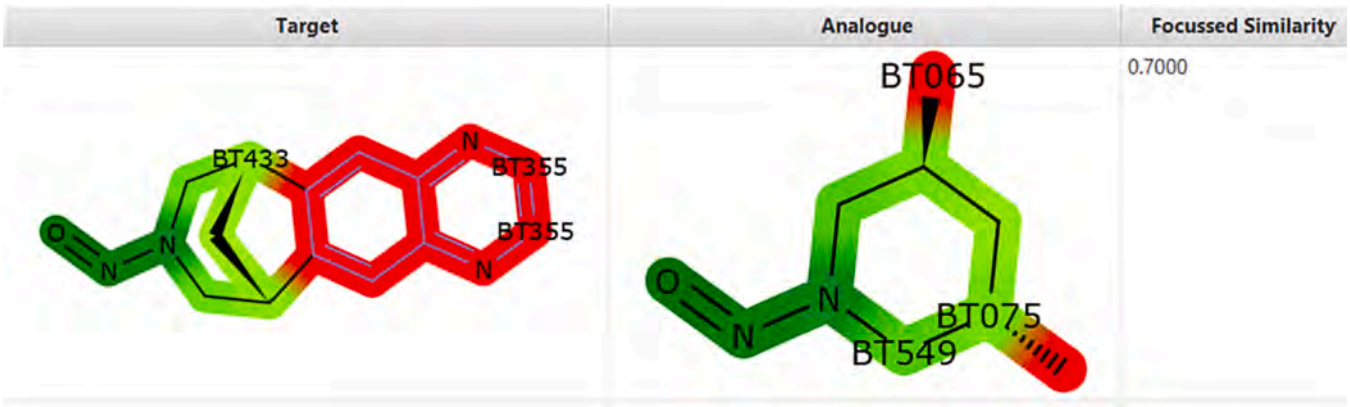

Fig. 7. Comparison of the target and one of the highest ranked candidates showing the predicted sites of metabolism on both molecules.

Compound	Focussed Similarity	Categories	Ames test	Ames test guideline compliance	Gold TD50	Lhasa TD50	Average Mass	Clearance (log10 mL/min/kg)	Log P	Solubility
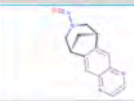	1.0000	Has_a_CH2 Large_side_chains Nitrosamine piperidine ring_size_6		Unknown.			240.26	Prediction: 0.869 95% confidence interval: 0.154 to 1.58	1.68	-2.66
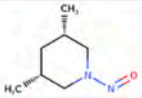	0.7000	Has_a_CH2 Nitrosamine piperidine ring_size_6	Positive	Strong. Has positive strain.			142.20	Prediction: 0.859 95% confidence interval: 0.143 to 1.58	1.92	-1.90
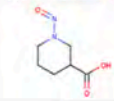	0.6706	Has_a_CH2 Nitrosamine piperidine ring_size_6	Negative	Weak. Missing strains: TA100, TA102, WP2 uvrA			158.16	Prediction: 0.427 95% confidence interval: -0.284 to 1.14	0.51	-0.30
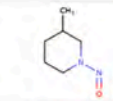	0.6667	Has_a_CH2 Nitrosamine piperidine ring_size_6	Positive	Strong. Has positive strain.			128.17	OUT OF DOMAIN. Prediction: 0.933 95% confidence interval: 0.217 to 1.65	1.68	-1.35
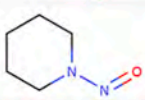	0.6386	Has_a_CH2 Nitrosamine piperidine ring_size_6	Positive	Strong. Has positive strain.	1.3, 1.43, 2.76, 12.1, 83.3	1.12, 29.1	114.15	OUT OF DOMAIN. Prediction: 0.880 95% confidence interval: 0.164 to 1.60	1.45	-0.95
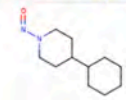	0.6000	Has_a_CH2 Nitrosamine piperidine ring_size_6	Negative	Medium. Four strains acceptable in most cases. +/- rodent S9 for all strains. Missing strain: TA102, WP2 uvrA			196.29	Prediction: 0.794 95% confidence interval: 0.0799 to 1.51	3.45	-3.35

Fig. 8. Extract from the prototype application summary table for the analysis of N-nitroso varenicline.

recommended to detect mutagens has been discussed [75] but its evaluation is outside the scope of this work, nevertheless, the guideline compliance is particularly important when considering data which suggest a compound is negative in the Ames test.

The prototype application provides the user with several graphical means to compare the similarity of the potential analogues with the target. Three different views are shown in Fig. 9: stacked bar chart, coxcomb plot and skyline plot; for the three views each of the selected properties is represented by a different colour. For the stacked bar chart, the width of the colour indicates the relative similarity of the property to that of the target compound. All the properties are scaled by the difference between the value for the analogue and that of the target with the greater width indicating more similarity. In the coxcomb plot, a circle is divided into segments of equal angle, with one segment per property and the similarity of the property of the analogue with that of the target represented by the length of the radius of the coloured part of the segment; the most similar analogues would fill the coxcomb plot circle. In the skyline plot the difference in the values of the selected properties for the analogues is shown normalised relative to that of the target but this view has the extra property of showing the direction of variation of the analogue's property relative to that of the target's property: positive differences being above the skyline and negative differences below the skyline – the focussed similarity will always be on or below the line as it cannot be greater than that of the target compound. For the skyline plot, the most similar analogues would show no differences, i.e. a flat skyline.

In Fig. 9 the selected analogues can be ordered either by the focussed similarity (A) or by the combined value of all the selected properties (B). In (B) the mass, calculated solubility values and calculated clearance values cause different analogues to rank higher than those with greater focussed similarity. Therefore, if it were believed that these properties were significant factors in the toxicity then the user might be inclined to consider the candidate analogues more or less suitable on this basis. The depiction of the similarity of the properties might be considered overly simple because there is no account taken of the error of the calculated properties: indeed the 95 % confidence interval for the clearance model calculations suggest that the values for some of the analogues are the same as that for the target. Furthermore, there is no weighting of the difference in value between the analogue and target for the different properties. Nevertheless, the display of properties in this way gives the user an indication of wider similarity considerations.

The charts in Fig. 9 show comparison of each analogue with the target; the prototype application also provides the user with charts that show comparisons for several compounds at once.

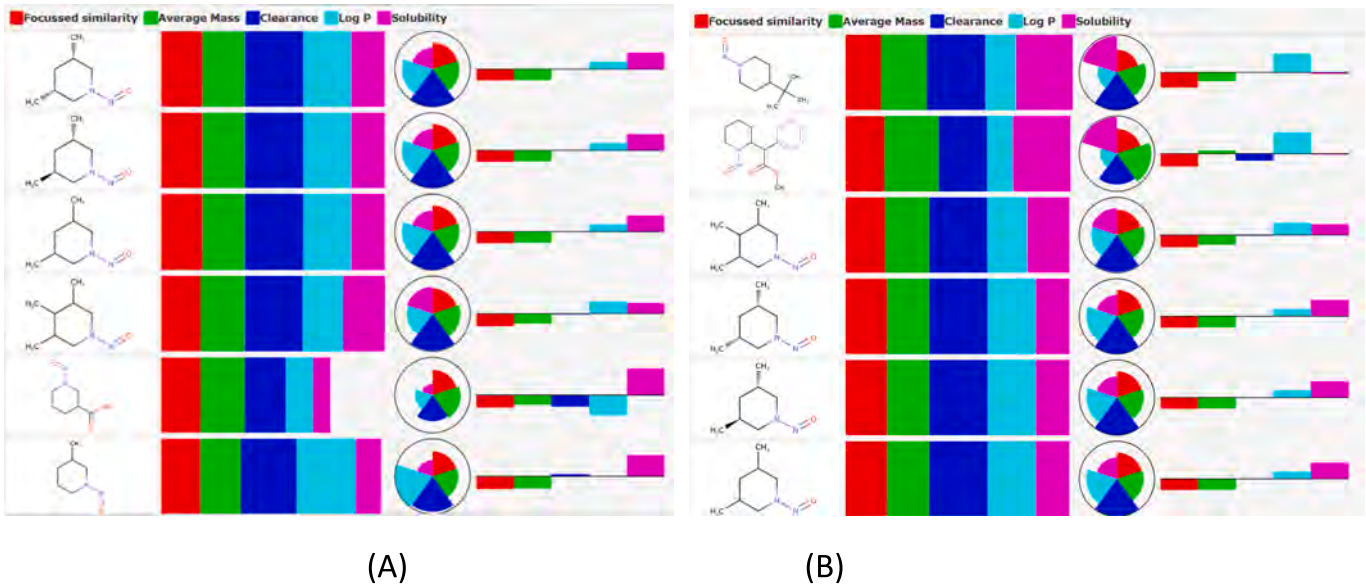

Fig. 9. Stacked bar chart, coxcomb plot and skyline plots showing focussed similarity and physico-chemical and biochemical measures for assessing the appropriateness of candidate analogues of N-nitroso varenicline. In (A) the analogues have been ranked by focussed similarity whereas in (B) they are ranked by taking into account all the displayed properties. In each of (A) and (B) the stacked bar char is on the left; the coxcomb plot is in the middle and the skyline plot is on the right. The selected properties and their colour coding is shown in the legend at the top of each of (A) and (B).

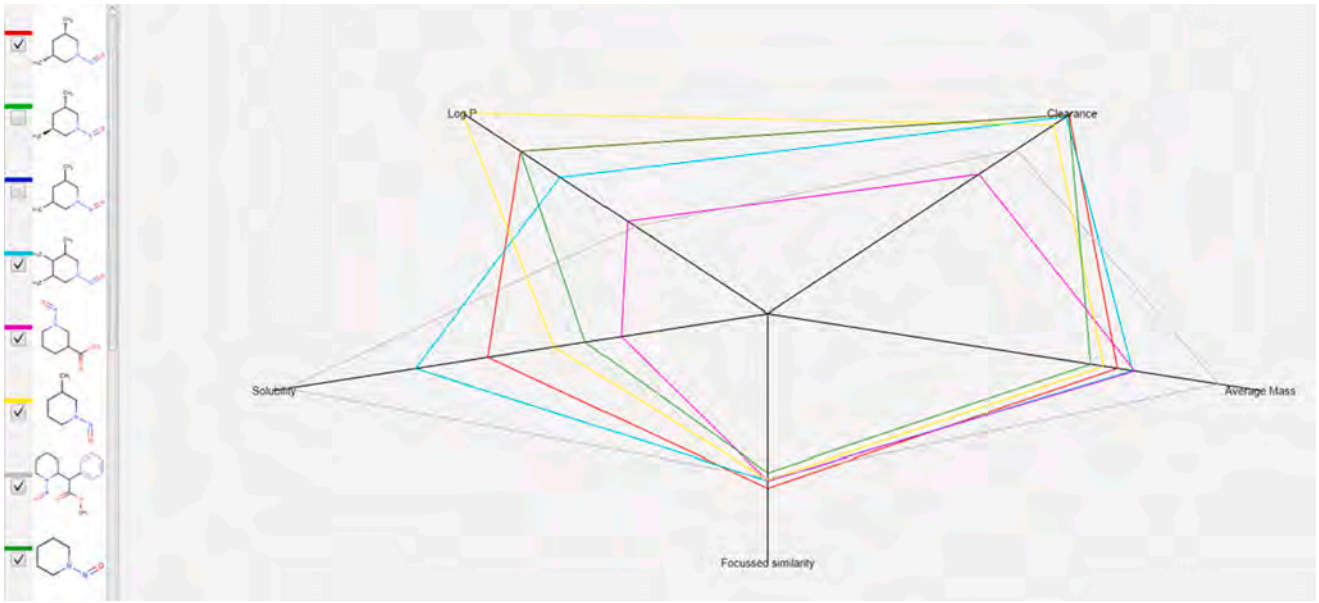

Fig. 10. Radar plot of selected potential analogues and their properties relative to the target compound. Longer distances represent increased similarity.

In Fig. 10, properties of some of the potential analogues are shown on a radar plot. Each axis represents one of the properties used to consider the appropriateness of the analogue and each of the selected analogues is represented on the plot by, in this case, a pentagon. In contrast to Fig. 9 the colours in Fig. 10 represent different analogues. Commonly, radar plots are created so that the shorter distance along each axis represents similarity, i.e. the item closest to the centre of the chart is the most similar to the reference, but in the case of Fig. 10 this is reversed so that the larger the shape, the more similar is the analogue to the target; this is a setting that the user can toggle, but the 'longer distance is more similar' configuration stops very similar analogues disappearing into a single point.

In Fig. 11, a selection of analogues and the similarity of their properties to that of the target are shown in a heat map, with blue representing values higher than the target and red representing values lower than the target; note that the focussed similarity will always be red or white as it cannot be greater than that of the target compound. Thus the display in Fig. 11 has the advantages of the skyline plot in Fig. 9 of allowing the user to see the direction of differences in the properties of analogues and target.

The charts in Fig. 9 – Fig. 11 are designed to complement each other and it is not expected that the user will use only one when selecting an appropriate analogue or set of analogues.

3.2.1.1. Conclusion for N-nitroso varenicline. Most of the candidate analogue compounds for *N*-nitroso varenicline are Ames positive and there are also data from rodent carcinogenicity studies for some of the analogues considered less similar. For the most similar compound, 3,5-

Fig. 11. Heat map representation of similarities of properties of selected analogues relative to the target compound. Red represents values lower than the target and blue represents values higher than the target. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

dimethyl, *N*-nitroso piperidine, Vitic has records for both this compound and its stereospecific *cis*- and *trans*- isomers which are shown to be Ames positive in strain TA1535 with rat S9 extract [76]. From this the user would probably conclude that the target compound will be Ames positive and carcinogenic.

The closest structural analogues support assigning the target a medium potency because they do not belong to a category related to low or high potency. This is comparable to recent regulatory guidance updates, where it has been assigned to CPCA category 3 and assigned an acceptable intake of 400 ng/day [77,78].

3.2.2. *Example 2: N-nitroso valsartan*

N-nitroso valsartan is an impurity that has been found in the angiotensin II receptor blocker valsartan [79]; the structures of these compounds are shown in Fig. 12. In the *N*-nitrosamine compound a *N*-nitroso group replaces an oxopentyl group. *N*-Nitroso valsartan has been found to be Ames negative [80] and an estimation of its toxicity using Derek and Sarah has been published [81].

Using the prototype application, firstly the data recorded for *N*-nitroso valsartan that are in Vitic can be examined; the data consist of negative results in five strains (*S.* Typhimurium TA98, TA100, TA 1535 and TA1537 and *E. coli* WP2 uva) both with and without S9 activation. This would probably be sufficient to allow the user to come to a conclusion without further investigation that the target is potentially not a mutagen and not a carcinogen via this mechanism. However, for the purpose of demonstration, and confirmation that the negative Ames results fit with similar compounds, a read-across was performed using the prototype application.

The nitrosamine group in *N*-nitroso valsartan is identified from the match on the Derek Nexus alert 070, relating nitro and nitroso compounds to carcinogenicity. *N*-Nitroso valsartan is found to be contained in seven different categories, *viz.* Has an alpha CH2; Has i-propyl; Acyclic; Benzylic; generic nitrosamine; Has a carboxylic acid; Has a beta methyl. Being in seven categories highlights one of the main drawbacks of the Venn diagram: only five categories can be displayed at any one time. Venn diagrams with more than five categories do exist but are not implemented in the prototype application.

A different view of the data is available within the prototype application, namely the UpSet plot which for the categories containing *N*-nitroso valsartan is shown in Fig. 13.

In Fig. 13, each category that *N*-nitroso valsartan falls into is listed in the lower part of the left-hand pane, shown as (A) in the Figure, with the number of possible analogues in the database belonging to that category displayed as a bar chart to the left. For example, there are 338 compounds in the database that fall into the has an alpha CH_2 (Has_a_CH2) category, whilst only fifteen fall into the has a beta methyl (Has_b_-methyl) category. The right-hand side of the plot, (B) in the Figure, shows the different combinations of the categories which contain one or more possible analogues, identified by the combinations of circles. Above each combination is the number of analogues in the database which fall into that combination of categories, shown as a bar chart (C). There are five compounds falling into three different combinations of categories indicated by (D) ; the five compounds are shown in Fig. 14. The UpSet plot in Fig. 13 contains, in the top of the left-hand pane, (E) an aide memoir of the potency ranges for each category, which also acts as a legend indicating the meaning of the colour of the category text in (A) and circles in (B).

Fig. 13 also shows the potency of each category by the colour of the category text in (A), and of each combination of categories by the colour of the circles identifying the combination in (B). For example, the combination of all categories, in (D), which contains only the target compound, is comprised of categories listed in both the low and high potency entries in Table 1;therefore this combination of categories is considered as being of medium potency though needing expert assessment.In contrast, the combination of categories in (D) that contains three compounds is comprised of two categories in Table 1 – carboxylic acid and has an iso propyl (Has_i_propyl) – considered to be potency lowering and so is considered to be of lower potency. Note the three compounds in this category combination may belong to categories which are not relevant to *N*-nitroso varenicline though are in Table 1 and so it is not true to say that all (or any) of the three compounds are considered low potency.

The detoxifying effect of the carboxylic acid group in mutagens and carcinogens is well known [82], and the carboxylic acid category

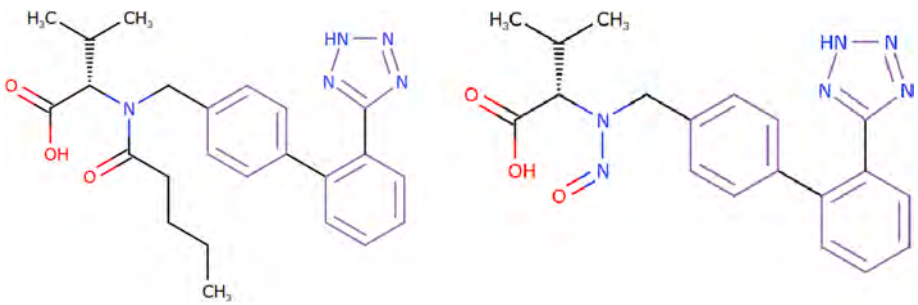

Fig. 12. Structures of valsartan (left) and N-nitroso valsartan (right).

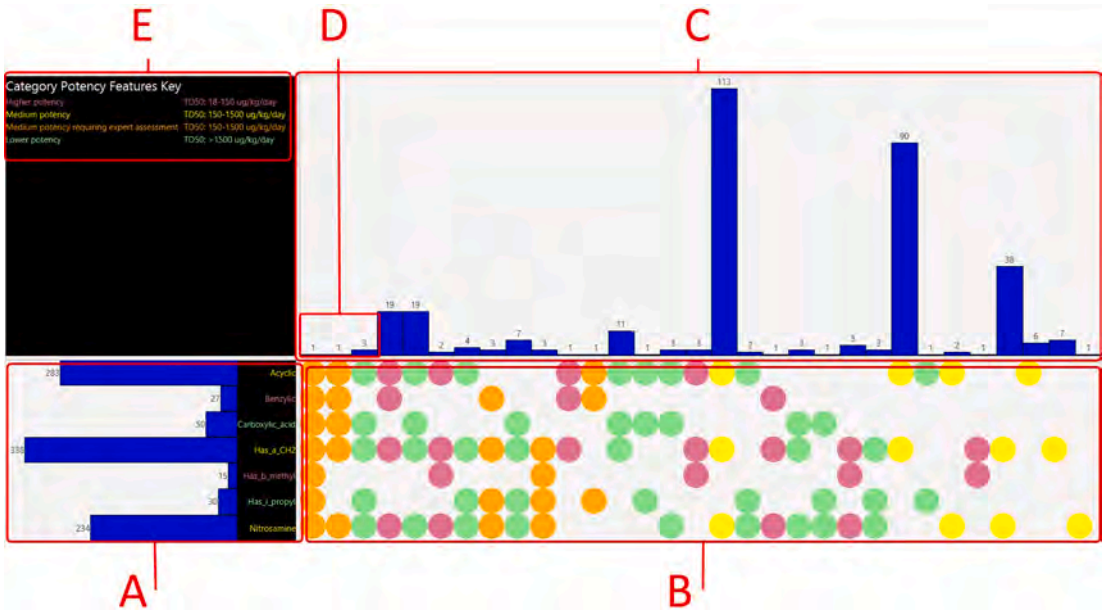

Fig. 13. UpSet plot of category combinations and populations for categories which contain N-nitroso valsartan. (A) List of categories in which the target compound falls along with a bar chart of number of possible analogues in each category. (B) Different combinations of categories. A circle indicates the inclusion of a category in the combination represented by each column. The colour of the text in (A) and the circles in (B) indicate the potency associated with the categories or their combinations. (C) Bar chart of the number of possible analogues in each combination of categories. (D) The combinations of categories, containing five compounds in total, which are explored in the text. (E) Key to the potencies of each colour of category combinations in (B).

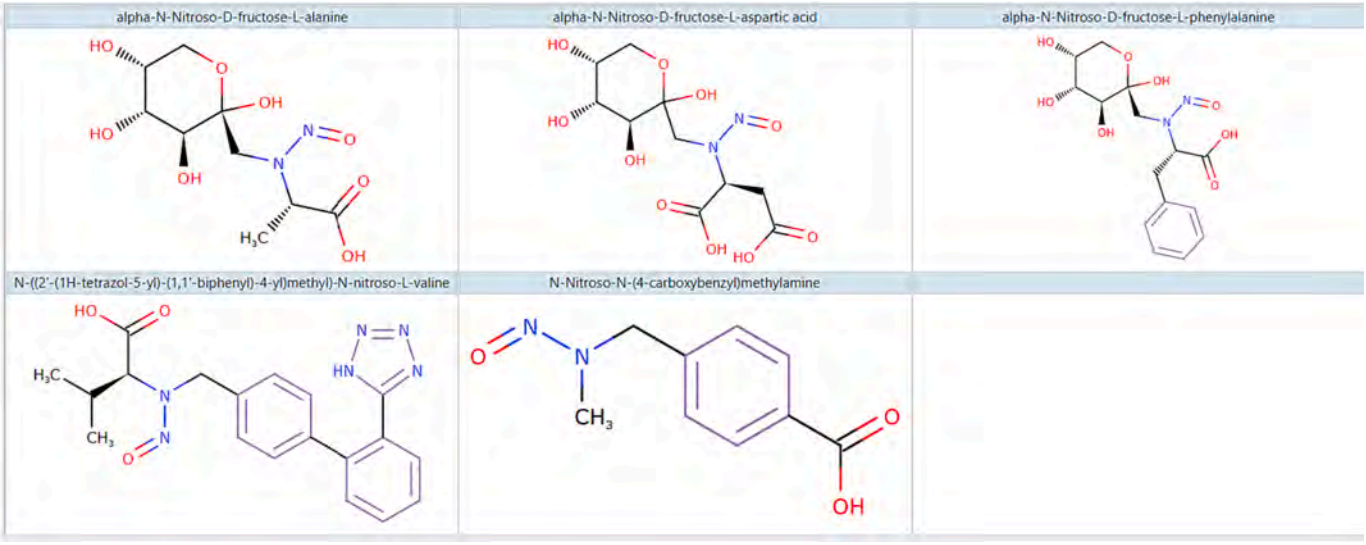

Fig. 14. Five compounds selected from the UpSet plot of combinations of categories containing N-nitroso valsartan.

definition allows the carboxylic acid to be anywhere in the molecule, however the user might wish to limit this. The prototype application allows the user to define their own categories by selecting a substructure in the target molecule. This is clearly advantageous in situations other than *N*-nitrosamines where categories are not defined. In Fig. 15 the user has selected the nitrosamine group and the carboxylic acid group of *N*-nitroso valsartan to define a category named beta Carboxylic acid. The prototype application allows the user to add restrictions to the selected atoms using the pattern language developed by Lhasa; for example in Fig. 15 the carboxylic acid oxygen has a hydrogen count restriction to ensure the category does not include carboxylic esters or peroxides etc.

User defined categories can be used alongside pre-existing categories, for example in the Venn diagram in Fig. 16 the user has selected their own category, beta Carboxylic acid and the pre-existing category of acyclic nitrosamines. The user may consider that the twelve compounds in the overlap of these categories represent better potential analogues than those chosen from overlaps of pre-existing categories alone.

The four compounds shown in Fig. 14 which are not the target are taken forward for further analysis, the first step of which is that they are ordered by focussed similarity. The focus of the similarity is the atoms and bonds that are defined by the toxicophore identified by Derek Nexus alert 070 and the similarity is calculated on the fully hydrogen-expressed structure; in this case a distance of six bonds gives the most informative separation of the different potential analogues and the substructure for the focussed similarity is shown in Fig. 17.

With these settings the four candidate analogues are ranked as shown in Fig. 18, which again shows areas of similarity between the target and the analogue coloured green for similar and red for dissimilar.

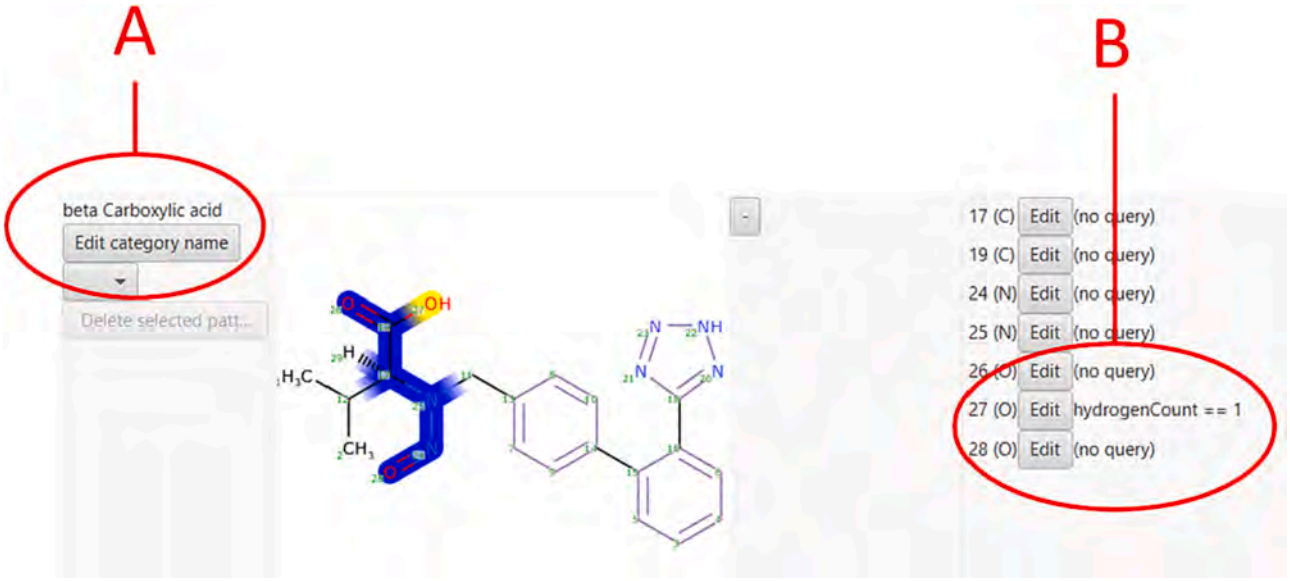

Fig. 15. User-defined category for beta carboxylic acids with atoms selected from the target compound highlighted in blue. The category name is entered at A. B shows the Explora pattern restriction on the oxygen, yellow highlighted in the structure, to ensure the category covers only acids and not esters etc. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

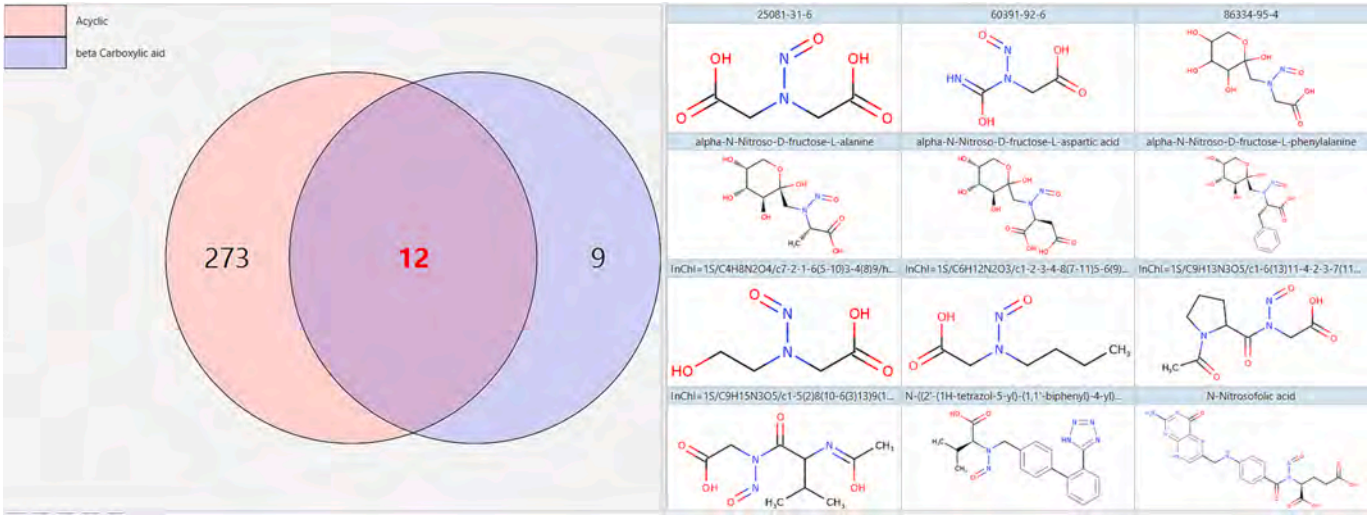

Fig. 16. Venn diagram of compounds contained in a user-defined category (beta Carboxylic acid) and pre-existing category (Acyclic). The overlap of these categories contains twelve compounds which are shown on the right.

Data for selected analogues can be investigated as shown in Fig. 19 with the same colour coding of the guideline compliance for the Ames data that was seen in Fig. 8. Note that in Fig. 19 three of the candidate analogues have reported negative Ames results in *S.* Typhimurium strain TA1538 which is not considered standard by the guidance.

The results of the investigation are summarised in table form alongside relevant toxicity data for the candidate analogues as shown in Fig. 20 and in a more visual form in Fig. 21. As can be seen in Fig. 20, all the candidate analogues are negative in the Ames test and on this basis, we would conclude that negative Ames results for *N*-nitroso valsartan are not anomalous and therefore it is also potentially not carcinogenic. The relevance of Ames data for predicting lack of carcinogenic activity in nitrosamines in general [83], and in sartans in particular, has been demonstrated [84].

When the process leads the user to come to a conclusion of no concern, as in this case, it is nevertheless important to consider factors such as the clearance of the target and how that is related to clearance of the analogue compounds. The large blue segments of the stacked bar chart and coxcomb diagram in Fig. 21 indicate that the predicted clearances of the analogues are similar to the predicted value of the target compound and thus remain suitable candidates for read-across to this target. In Fig. 21 an additional calculated property – topological polar surface area, one of many calculated properties available to the user – has been included; its variation over the possible analogues contrasts with that of log P and solubility. The figure shows how perfect correlation appears in the bar chart, coxcomb and skyline plot because *N*-nitroso valsartan is included in the set of compounds under consideration. The Figure also shows how the candidate analogue considered third most similar by focussed similarity is considered most similar overall when the physico-chemical and clearance parameters are taken into account.

3.2.2.1. Conclusion for N-nitroso valsartan. Even though there are data for the target compound itself, i.e. that it is Ames negative, it is

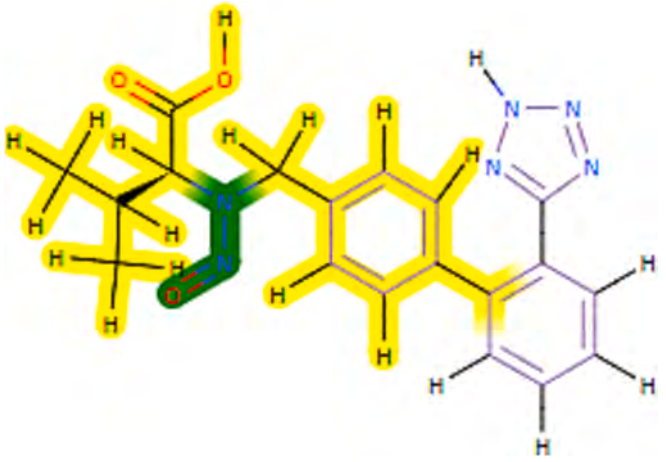

Fig. 17. Toxicophore highlighted on the fully hydrogen-expressed structure of N-nitroso valsartan.

instructive for the user to consider that there are many features related to the *N*-nitrosamine which makes it covered by seven different categories. This in turn places the target compound as needing expert review. All the closest analogues considered in the example are also Ames negative, but it is legitimate to consider if the potential analogues are close enough to the target to be persuasive: that *N*-nitroso valsartan falls into so many different categories is an indication of the complexity of its structure in relation to *N*-nitrosamine toxicity. In comparison, the CPCA algorithm classifies *N*-nitroso valsartan as category 4 with an AI of 1500 ng/day [68].

3.3. Validation

It is difficult to have a definitive approach to validating a read-across workflow as the process is expert driven with many decisions that are made having a certain subjectivity to them [33]. However, in order to investigate how much information may be derived algorithmically, rather than with expert input, from the different similarity approaches encoded in the prototype application we considered two aspects: (i) the use of read-across to make a correct call on the activity (positive or negative) of a *N*-nitrosamine and (ii) the use of read-across to estimate the potency, expressed as TD_{50}, of those nitrosamines where a TD_{50} was known.

It is worth re-emphasising that this validation study, being algorithmic, is expected to come to less convincing and less reliable conclusions than would a human expert using the prototype application manually. The case of *N*-nitroso valsartan described above is a clear example of this where a human user would select analogues within the carboxylic acid category; the algorithmic approach described here does not discriminate between the categories in this way. A more sophisticated algorithmic approach than the one employed here, and which might better reflect the assumptions and knowledge of an expert user, could use the predictivity of each of the categories – i.e. how well a category is related to either positive or negative classification – and weight likelihood of correctness accordingly.

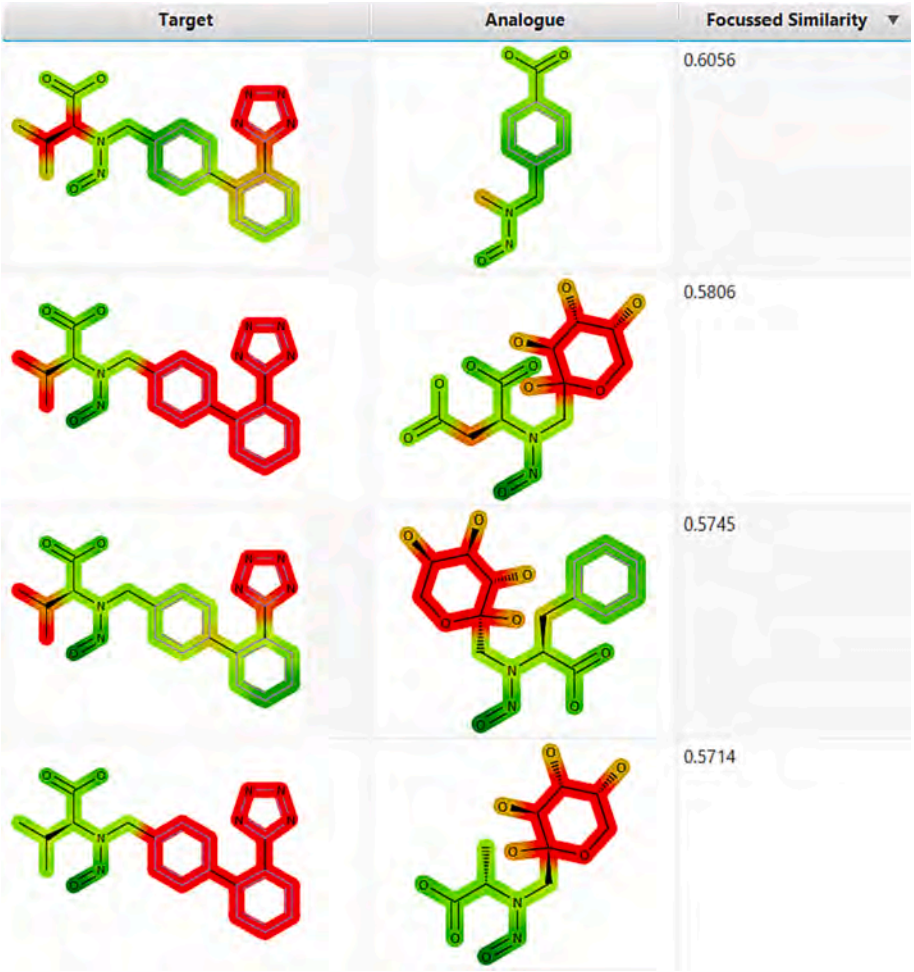

Fig. 18. Areas of similarity and focussed similarity score of four candidate analogues of N-nitroso valsartan.

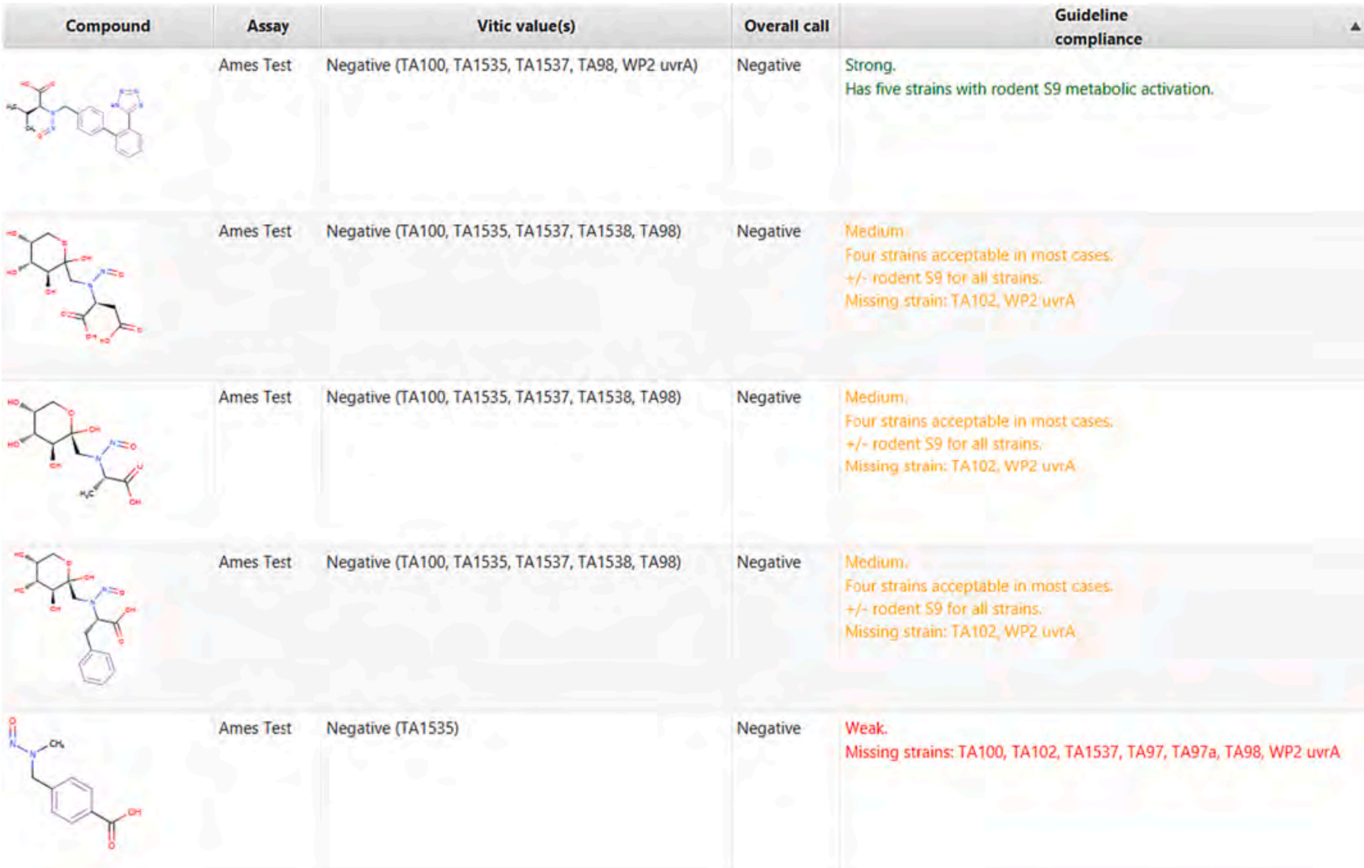

Compound	Assay	Vitic value(s)	Overall call	Guideline compliance
	Ames Test	Negative (TA100, TA1535, TA1537, TA98, WP2 uvrA)	Negative	Strong. Has five strains with rodent S9 metabolic activation.
	Ames Test	Negative (TA100, TA1535, TA1537, TA1538, TA98)	Negative	Medium. Four strains acceptable in most cases. +/- rodent S9 for all strains. Missing strain: TA102, WP2 uvrA
	Ames Test	Negative (TA100, TA1535, TA1537, TA1538, TA98)	Negative	Medium. Four strains acceptable in most cases. +/- rodent S9 for all strains. Missing strain: TA102, WP2 uvrA
	Ames Test	Negative (TA100, TA1535, TA1537, TA1538, TA98)	Negative	Medium. Four strains acceptable in most cases. +/- rodent S9 for all strains. Missing strain: TA102, WP2 uvrA
	Ames Test	Negative (TA1535)	Negative	Weak. Missing strains: TA100, TA102, TA1537, TA97, TA97a, TA98, WP2 uvrA

Fig. 19. Data retrieved from Vitic for N-nitroso valsartan and some potential analogues showing the suggested compliance with guidelines for regulatory acceptance.

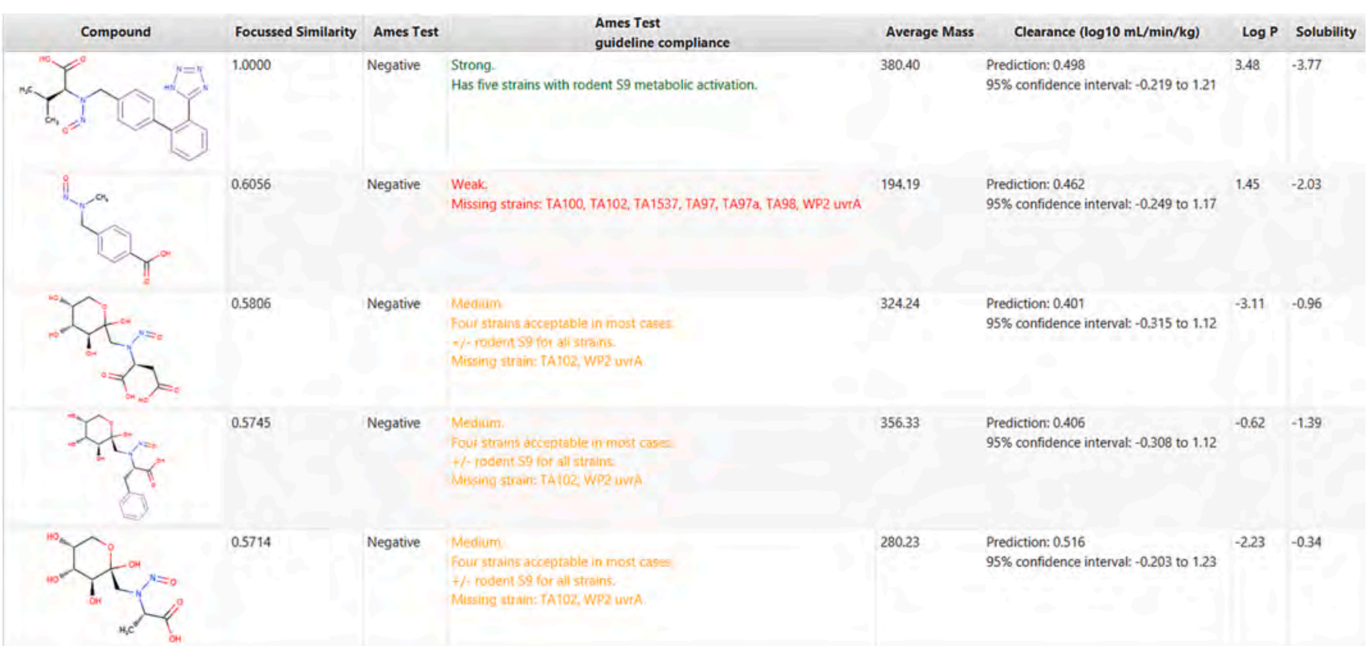

Compound	Focussed Similarity	Ames Test	Ames Test guideline compliance	Average Mass	Clearance (log10 mL/min/kg)	Log P	Solubility
	1.0000	Negative	Strong. Has five strains with rodent S9 metabolic activation.	380.40	Prediction: 0.498 95% confidence interval: -0.219 to 1.21	3.48	-3.77
	0.6056	Negative	Weak. Missing strains: TA100, TA102, TA1537, TA97, TA97a, TA98, WP2 uvrA	194.19	Prediction: 0.462 95% confidence interval: -0.249 to 1.17	1.45	-2.03
	0.5806	Negative	Medium. Four strains acceptable in most cases. +/- rodent S9 for all strains. Missing strain: TA102, WP2 uvrA	324.24	Prediction: 0.401 95% confidence interval: -0.315 to 1.12	-3.11	-0.96
	0.5745	Negative	Medium. Four strains acceptable in most cases. +/- rodent S9 for all strains. Missing strain: TA102, WP2 uvrA	356.33	Prediction: 0.406 95% confidence interval: -0.308 to 1.12	-0.62	-1.39
	0.5714	Negative	Medium. Four strains acceptable in most cases. +/- rodent S9 for all strains. Missing strain: TA102, WP2 uvrA	280.23	Prediction: 0.516 95% confidence interval: -0.203 to 1.23	-2.23	-0.34

Fig. 20. Tabulated results for N-nitroso valsartan and candidate analogues.

3.4. Activity call

We took a set of 494 nitrosamines with data relevant to the carcinogenicity endpoint (418 positives and 76 negatives) and used three methods to make a call as to their activity. Firstly, taking the analogues which were in different numbers of categories shared with the target; secondly taking the most similar compound(s) by focussed similarity and thirdly combining these approaches by taking the potential analogues with the most categories in common with the target and ranking them by the focussed similarity. This last method is most close to the

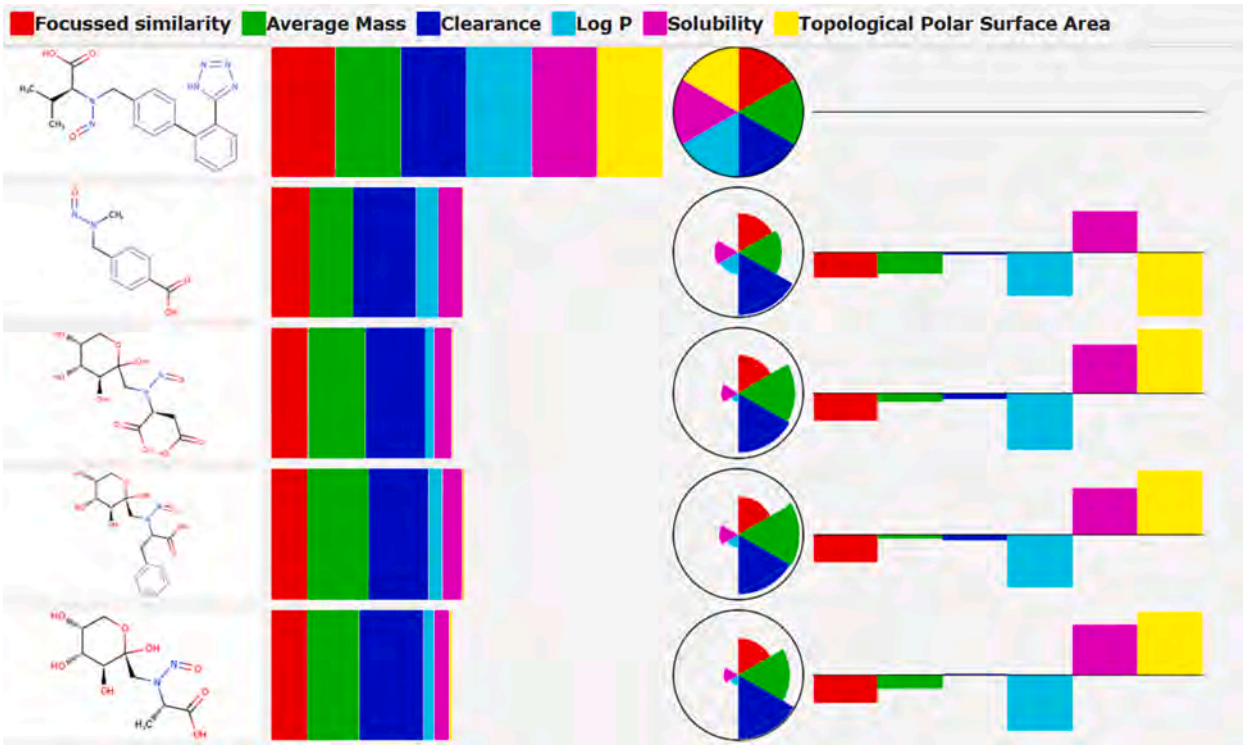

Fig. 21. Graphical summary of similarities N-nitroso valsartan and four analogues.

workflow described in the examples earlier in the paper. The three methods were applied to each of the compounds in the dataset in a leave-one-out cross validation approach.

With such a biased dataset, a positive call at random would very likely be correct and so the challenge is to see how well the negative compounds can be identified (the true negative rate or specificity). To see if the algorithmic approaches were correctly classifying the members of the dataset in a way that was better than a random call, the results were compared to the distribution of correct predictions made by the same method but using a y-randomised dataset – that is a dataset in which the positive or negative labels were randomised – thus producing a dataset with the same positive and negative occurrence rate. The randomisation of the data was run 10 times and the true positive rate and true negative rate was recorded for each run.

The correctness of prediction, as measured by the true positive rate and the true negative rate, was correlated with the number of categories that each compound appeared in. This was considered in two ways: (i) the absolute number of categories that potential analogues shared with the target compound, Fig. 22 and (ii) the number relative to the highest number of categories that any potential analogue shared with the target

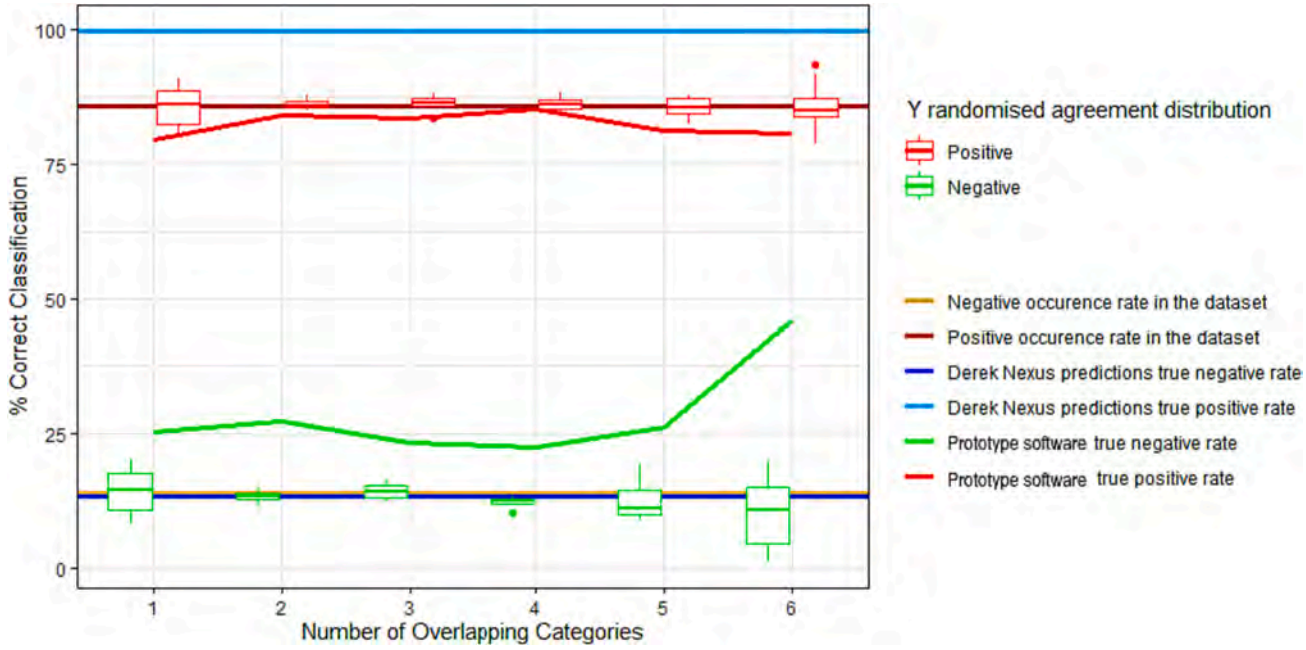

Fig. 22. Percentage correct classification of activity of a set of 494 nitrosamines relative to the number of categories analogues share with the target.

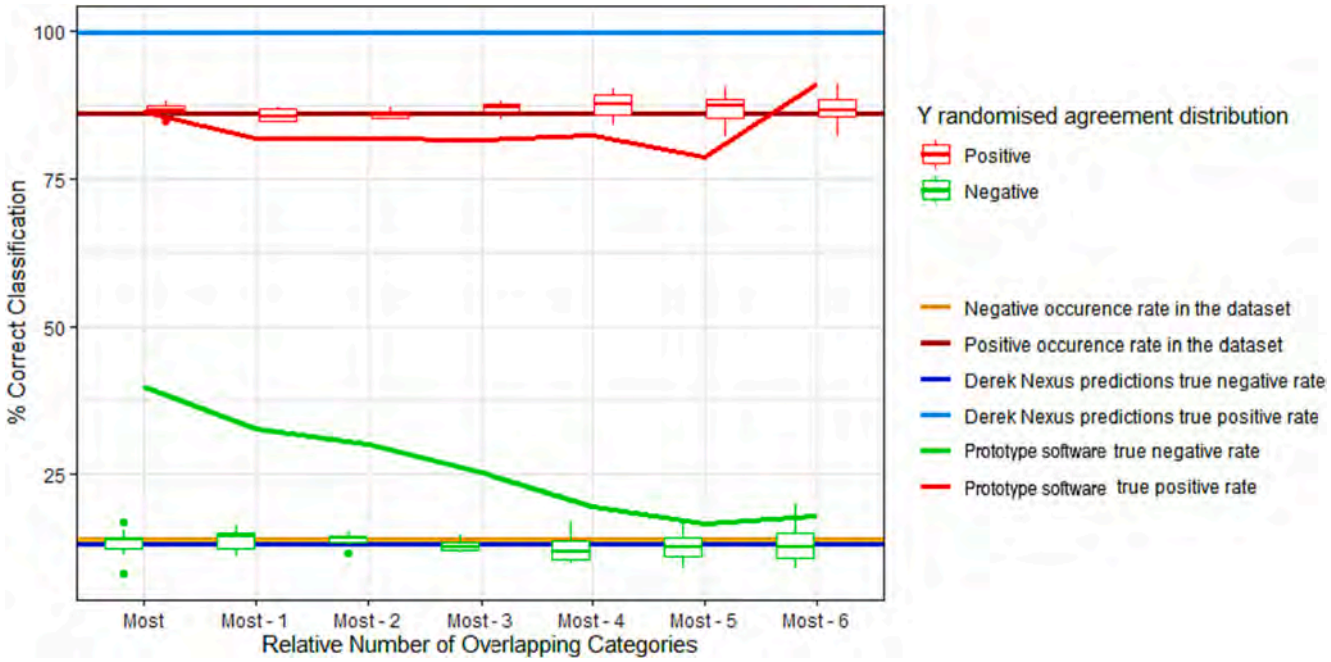

Fig. 23. Percentage correct classification of activity of a set of 494 nitrosamines relative to the relative number of categories analogues share with the target.

Fig. 23.

Both Fig. 22 and Fig. 23 show the variation of true positive rate and true negative rate varying with the number of categories that an analogue has in common with the target. The y-randomised data are shown in box plots in the same figure, as distributions of true positive and true negative rates. The two figures also include the rates of occurrence of positive and negative activity in the dataset along with the correct prediction rate of Derek Nexus for purposes of comparison.

As can be seen, in both cases the true positive rate is fairly independent of the number of categories that a potential analogue has in common with the target compound and is generally close to the occurrence ratio in the dataset (and somewhat less than that of the Derek Nexus predictions). From this we conclude that the algorithmic approach adds nothing when identifying positive *N*-nitrosamines. However, the true negative rate shows increasing correctness with increasing numbers of categories in common between the candidate analogues and the target. A smoother relationship appears when the number of categories in common between the analogue and the target are adjusted relative to the highest number for each target (Fig. 23).

In considering the usefulness of the classification of activity, as was mentioned earlier, it is probably more important to identify negatives among the many positives. Thus, in Fig. 23, correctly detecting about 40 % of the negatives represents a significant improvement but note that this is only achieved where analogues share a large number of overlapping categories with the target. Where potential analogues and the target share fewer categories, the algorithmic approach performs little better than random.

The second attempt to understand the value on the underlying data, namely using the focussed similarity, took the same dataset and again used a leave-one-out cross validation compared to a set of y-randomised cross validation runs. In this approach, the distance from the focus of similarity was measured incrementally until only one compound had the highest focussed similarity. A maximum of ten atoms distance was used. If no single compound were identified as having the uniquely best similarity to the target, then all the compounds having the highest similarity at the end of the process were considered equally similar to the target. The call of positive or negative was made on this activity of the single most similar compound (or an average of them where no single most similar compound could be identified). The useful true negative rate by this technique was poorer than that of the technique using the number of categories that the analogues had in common with the target, as shown in Fig. 24.

A third experiment combined these two approaches by taking the set of compounds with the highest number of categories in common with the target and using the focussed similarity to identify a most similar compound. The call of positive or negative was, as with the previous experiment, made on the activity of the single most similar compound. This produced a better true negative rate than using the focussed similarity alone, but not quite as good as using only the number of categories in common between the analogues and the target, again as shown in Fig. 24.

The dataset used for the above experiments contained a significant number of compounds whose activity call was based only on the Ames result. A dataset was therefore also constructed as a subset of the one used above in which those compounds which only had Ames data were excluded. This produced a set of 250 compounds (201 positive and 49 negative) in which the same experiments as above were run. Fairly similar results were obtained from these experiments. Fig. 25 to Fig. 27 show the results of these experiments. As can be seen, the clearer trend in relating number of overlapping categories to, in particular, true negative rate occurs when the number of categories is sorted relative to the highest rather than when the absolute number of categories is used. The results of the classification based on highest number of categories in common is very similar for this subset (40 % true negative rate, 86 % true positive rate) despite the slightly higher negative occurrence ratio in the set.

We can conclude from this study that the information about *N*-nitrosamines being considered negative is held more in the categories than in the focussed similarity, whereas the reverse is true for compounds considered positive. We can explain the success of the number of categories in common between the analogues and the target at predicting true negative rate over that of the focussed similarity methods because categories describing the presence of deactivating groups such as carboxylic acid will be considered together, whereas for focussed similarity metrics the deactivating carboxylic acid will be at different distances from the nitrosamine group and thus will only be considered if they are in an analogue that has the carboxylic acid at the same distance from the nitrosamine in the target.

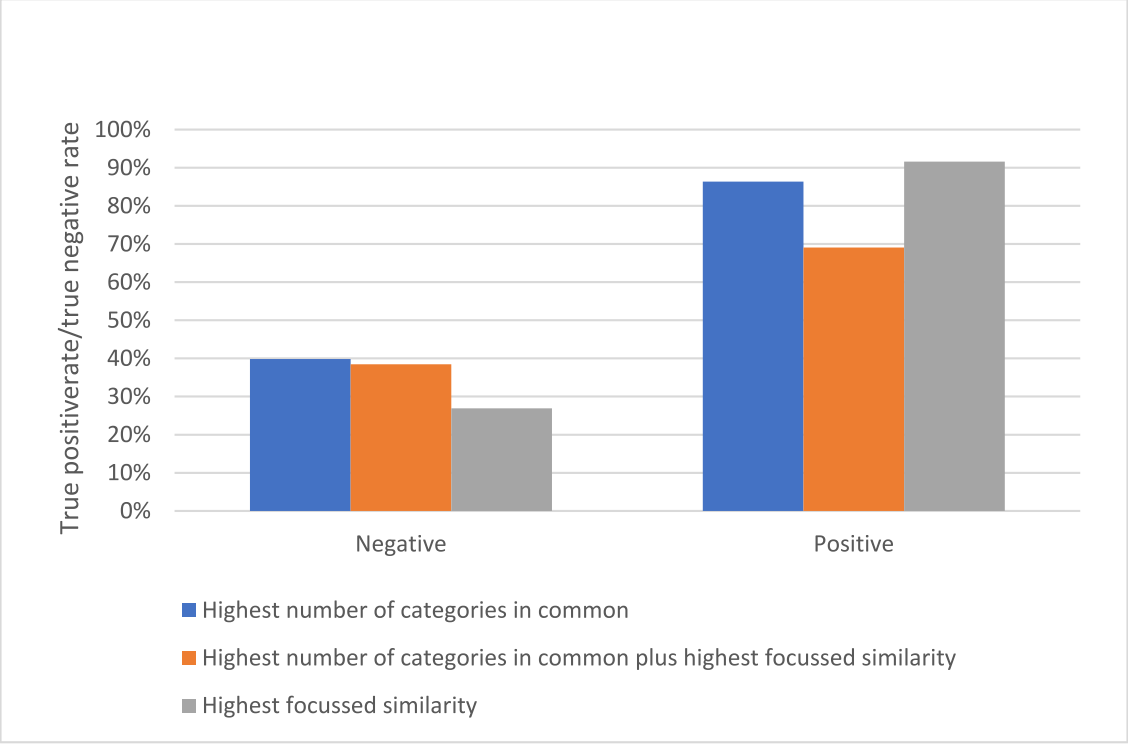

Fig. 24. Comparison of true positive rate and true negative rate for three different methods of selecting analogues.

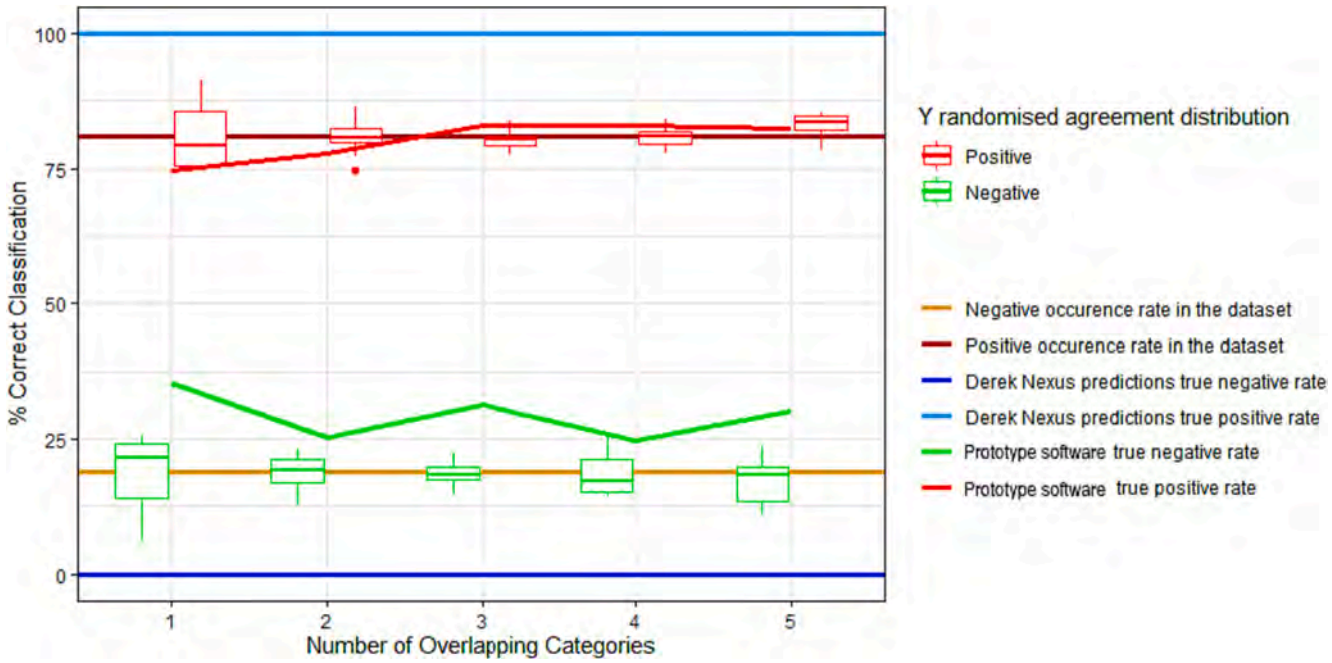

Fig. 25. Percentage correct classification of activity of a set of 250 nitrosamines with data other than Ames relative to the number of categories analogues share with the target.

3.5. Estimation of potency

The estimation of potency used the same three techniques for identifying analogues as described above, *viz.* number of categories in common between the analogue and the target. However, for this experiment, only *N*-nitrosamines with known TD_{50} values were used. This was a much smaller set of only 48 *N*-nitrosamines, which all appear in the LCDB. All the compounds had at least one recorded value of both the Lhasa TD_{50} and the Gold TD_{50} and for the experiment we used both measures though never combined them. Where there was more than one value recorded of either Lhasa TD_{50} or Gold TD_{50} – a situation that arises when separate TD_{50}s are recorded for different species and/or sexes – a mean value was taken; the data therefore represent a generic mammalian value for Lhasa TD_{50} or Gold TD_{50} rather than anything more specific.

With so few compounds, the results suffer from thirty compounds

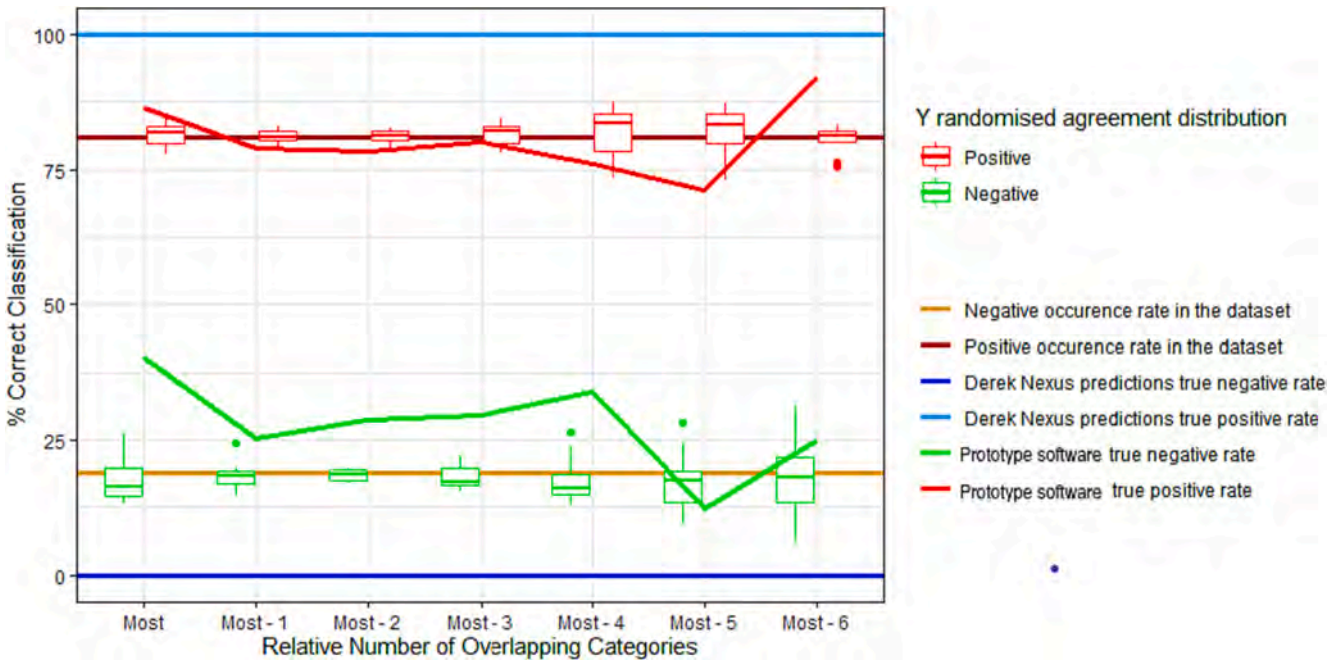

Fig. 26. Percentage correct classification of activity of a set of 250 nitrosamines with data other than Ames relative to the relative number of categories analogues share with the target.

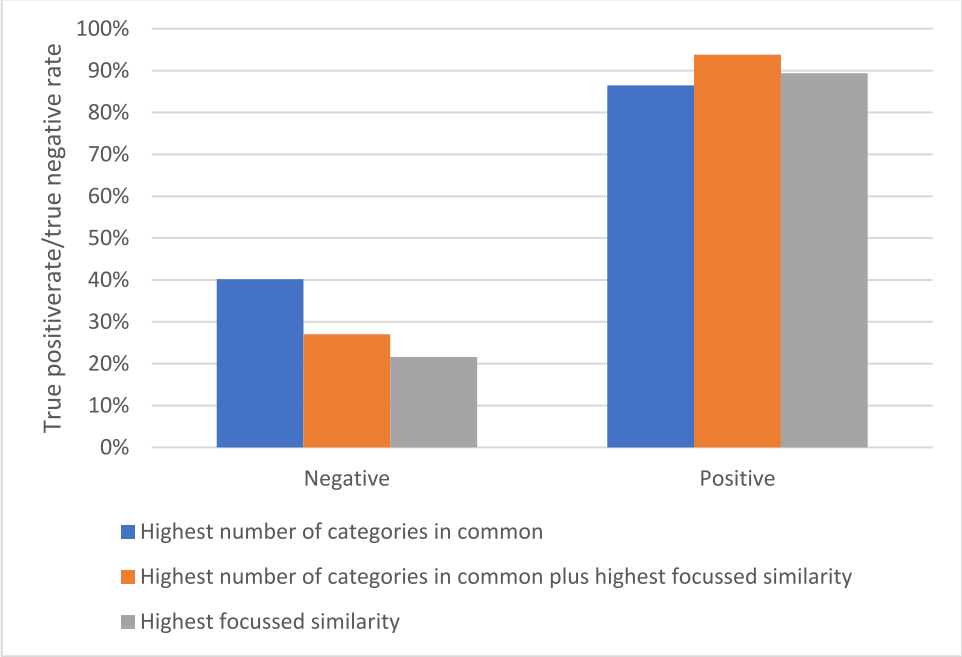

Fig. 27. Comparison of true positive rate and true negative rate for three different methods of selecting analogues for a set of 250 nitrosamines with data other than Ames.

falling into fifteen pairs where each finds the other the most similar and thus the error in the predicted TD_{50} is replicated in these pairs. The measure of error in the estimation was mean relative error (MRE) where the relative error is defined in Eq. (1). This is a good measure because it gives greater values for errors on small values (i.e. the more toxic compounds); a difference of 1 mg/kg/day makes little difference if the true TD_{50} is 100, but a huge difference if it is 0.01!

$$Relativeerror = \left| \frac{Calculatedvalue - Expectedvalue}{Expectedvalue} \right| \quad (1)$$

This penalising of errors on small observed values illustrates a limitation of the read-across approach for estimation of potency, namely that when doing a read-across from a member of a set of compounds there is a *de facto* applicability domain and that compounds whose value, TD_{50} in this case, lies outside the range of values contained in the set cannot effectively be estimated by read-across. This is particularly the case where the target compound is more toxic than any other member of the set. In our leave-one-out cross validation experiments this situation pertains when it is the most potent member of the set that becomes the target compound.

Fig. 28 shows the distribution of Lhasa TD_{50} and Gold TD_{50} values for the dataset; the values are shown on a log scale for ease of viewing, log values were not used in the validation experiments. As can be seen, there is one compound for which the Lhasa TD_{50} is more than an order of magnitude smaller than the rest of the values in the set. This compound, methyl, 2-oxo-1-propyl nitrosamine has several of the potency increasing features identified by Cross and Ponting [25]. Manifestly, if either an expert or an algorithm were to choose analogues for this compound from the set that is available, the relative error will be very large. For this reason, this compound was considered outside the applicability domain for the read-across experiment.

As with the experiments to investigate activity call, the MRE value was compared to a set of experiments in which the observed values were randomised over the dataset. The number of y-randomised experiments in this section was increased to 100 to take into account the smaller dataset size.

The most successful experiment used only the number of overlapping categories to select an analogue and the results are shown in Fig. 29 where it can be seen that the MREs in the estimated TD_{50} values fall below, or towards the lower end of, the range of those calculated at random. As might be expected, the increasing number of overlapping categories that an analogue has with the target generally leads to a better estimation of the TD_{50} and these are around 1 – 2 orders of magnitude where the number of categories is greater than three.

Of course, in all these validation experiments, the algorithmic approach will underperform that of a human expert examining the same data. But the results do suggest that the number of categories in common between analogue and target contains a lot of information about selecting an analogue for both negative activity and potency of *N*-nitrosamines.

4. Discussion

Selecting a read-across analogue is very often a process in which expert knowledge, judgement as well as some subjectivity are used. However, the factors relevant to making a read-across call can be identified and agreed on, even if their relative weight in the decision-making process is not agreed. In the case of *N*-nitrosamines, the pressing need is to make a call on whether or not the compound is carcinogenic and, if so, what its likely potency is. The carcinogenic/non-carcinogenic call is easier to make than the potency. The precision of the potency call – either reading across a point value or assigning the target compound to a range of potencies – is necessarily difficult given the lack of data. Having a software tool that supports gathering all the necessary data and thus keeping the workflow consistent as far as possible helps the user make consistent decisions. It is also important to be able to acknowledge where there is insufficient data or where a suitable analogue cannot be identified.

The workflow described above and its validation, in terms of the contribution of different ways of considering similarity, can be supported through implementation into user-friendly software. One aspect of the validation worthy of further consideration is the relative success of the categories in providing information about both the activity and potency of the nitrosamines studied, relative to the focussed similarity. It is valuable to add a little more detail to the discussion by considering the TD_{50}s dataset. The sets of both Lhasa TD_{50} and Gold TD_{50} values have very skewed ranges with most of the data points being closely clustered, as shown in Fig. 30; this makes the y-randomised values relatively successful at estimating each other. Nevertheless, the algorithmic read-across based on the number of categories that the analogue and target have in common consistently does better.

Ideally, the dataset would be large enough to validate how frequently the observed value for the target TD_{50} fits within the range of the analogues selected, in particular within the range specified by particular quantiles. However, the dataset is too small and too biased to make such an analysis informative. Nevertheless, the success of the using the number of categories as an estimator of the activity and potency of a *N*-nitrosamine can be rationalised. The set of forty-five *N*-nitrosamine categories can be considered to be the elements of a very expertly curated binary fingerprint, with forty-five bits. Analysis of the 494 *N*-nitrosamines used in the dataset suggests that the fingerprint is relatively sparsely populated with a distribution of cardinality as shown in Fig. 31, which also shows the cardinality of the pairwise intersection of category fingerprints for these compounds.

The validation exercise shows that overlap of the different categories provides a very effective means of identifying both *N*-nitrosamines

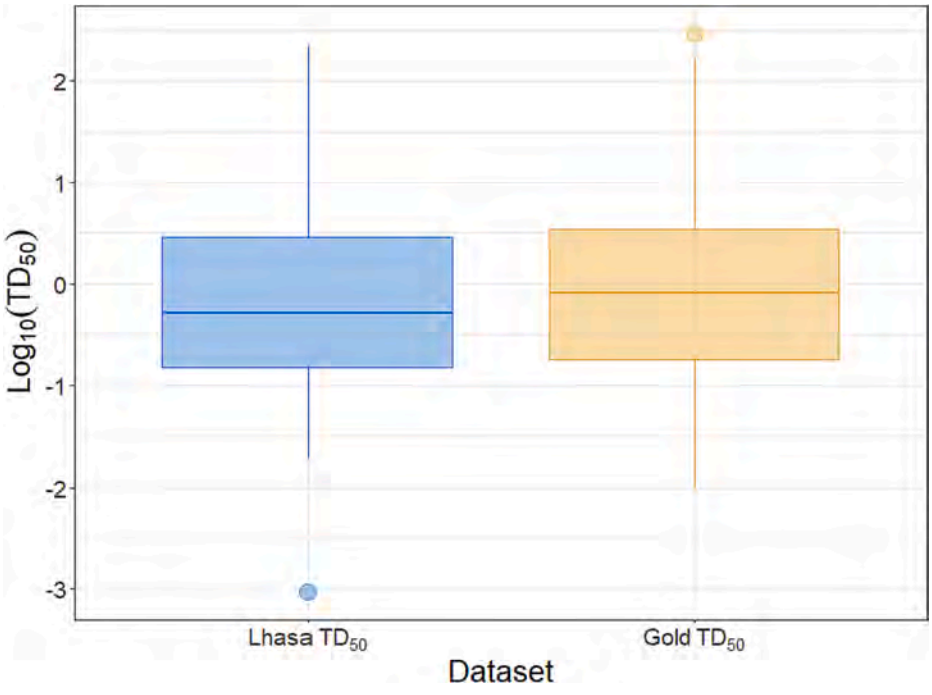

Fig. 28. Box plot showing the distribution of Lhasa TD_{50} and Gold TD_{50} values in the dataset used for the validation studies.

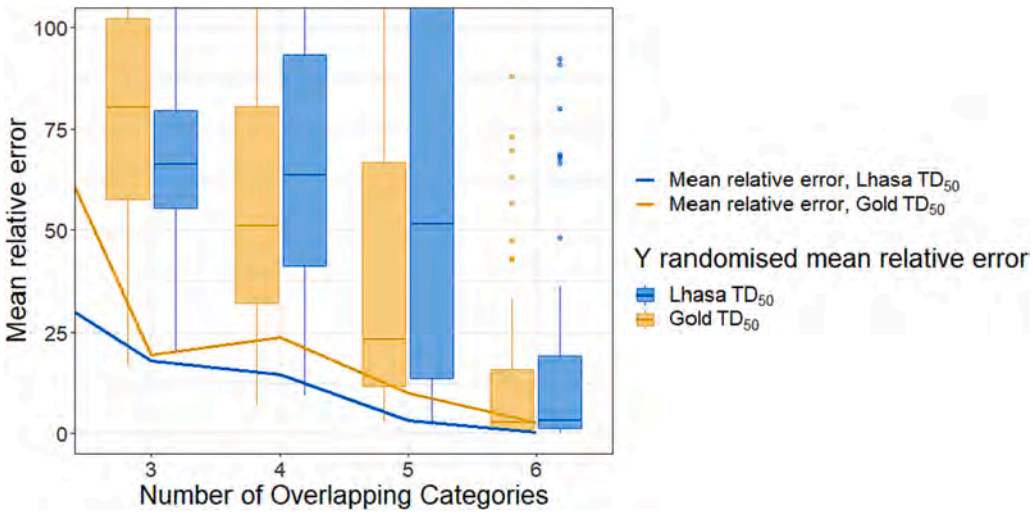

Fig. 29. Box plot of mean relative error of Lhasa TD_{50} and Gold TD_{50} estimation using varying number of overlapping categories compared to the distributions of same measures from a set of y-randomised experiments.

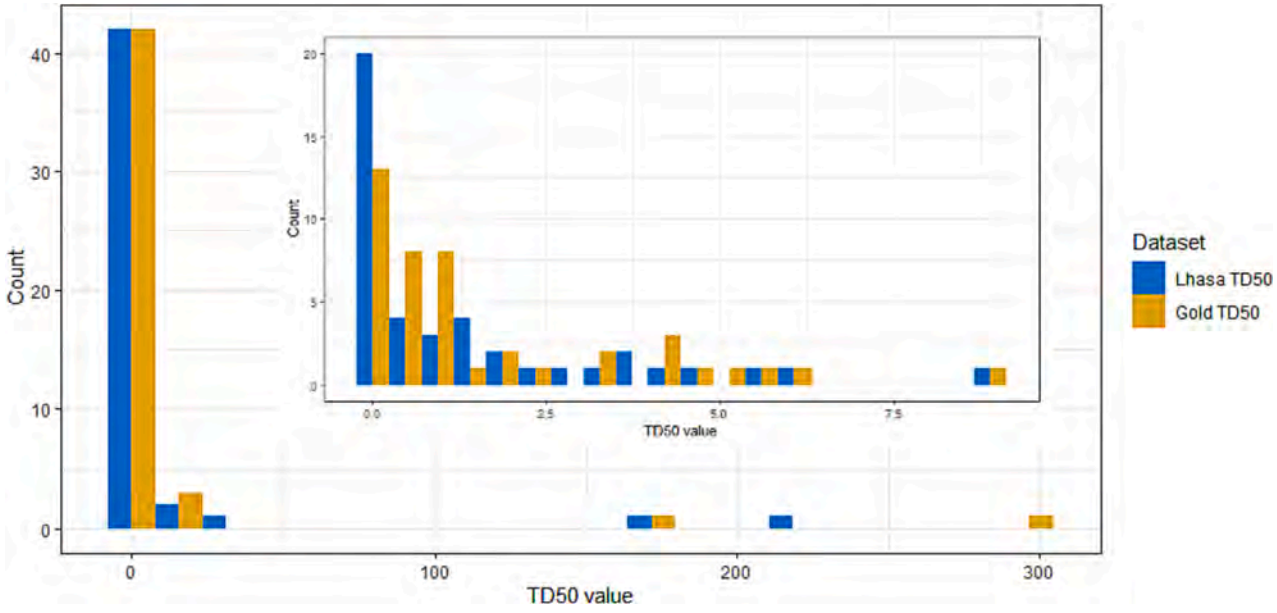

Fig. 30. Histogram of Lhasa TD_{50} and Gold TD_{50} values for a set of 48 N-nitrosamines and, inset, the lower value end of the range in detail.

considered not to be potentially carcinogenic and estimating the potency of those which are potentially carcinogenic. As was mentioned above, the identification of *N*-nitrosamines as negative carcinogens due to the presence of the carboxylic acid group is not very dependent on the position of the carboxylic acid relative to the *N*-nitrosamine and thus the category fingerprint deals better with this situation than does the focussed similarity. Furthermore, that the categories are derived by human experts means that they will capture the most important factors in influencing both activity and potency – thus a relatively small number of factors outperform the focussed similarity fingerprint consisting of many structural descriptors which may not be relevant to *N*-nitrosamine toxicity.

5. Conclusions

The use of the prototype application has demonstrated how software support for read-across can help the user consider all the different aspects of selecting a category or analogue in the analysis.

In the context of *N*-nitrosamines the use of expert-defined categories, and their consideration as part of the assessment helps the user by focussing on known factors in the making the assessment. Furthermore, using a defined toxicophore to focus a calculated similarity can produce potential analogues with a different emphasis. The use of the two techniques illustrates how the user may want to consider different aspects of the target molecule which are hard to capture *a priori*.

Inclusion of physicochemical and/or pharmacokinetic properties of the target compound presents yet another dimension to the selection of analogues, as does consideration of the metabolic degradation of the target in comparison to both potential analogues and other structurally related compounds (e.g. parent amines of nitrosamines).

Finally, the prototype application illustrates how the all-important data for analogues can be assessed in the context of all the factors above. The multiplicity of different factors that someone making a read-across assessment has to take into consideration emphasises how a software tool that supports these factors as part of a well-defined, though user-controlled, process can make the read-across analysis more

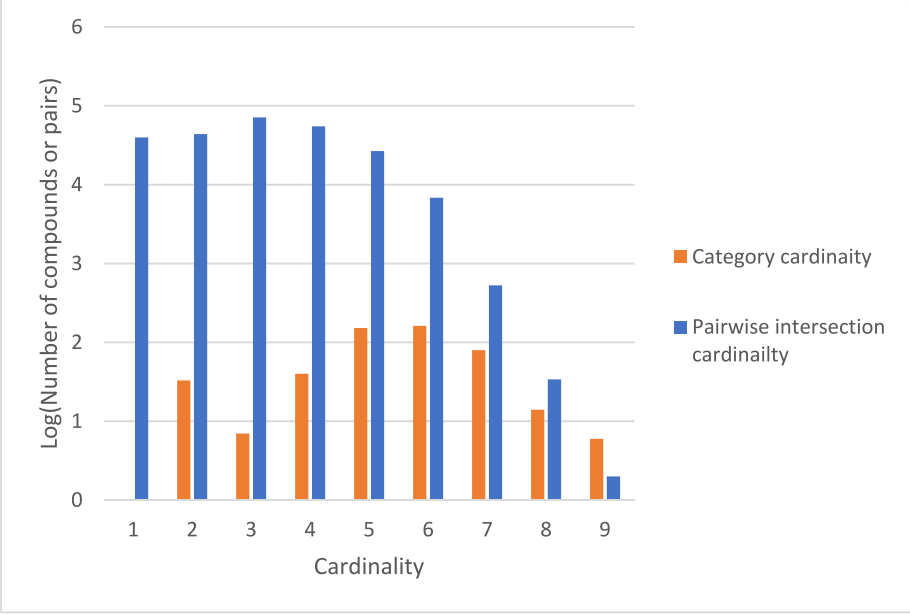

Fig. 31. Distribution of category cardinality and pairwise intersection cardinality for 494 N-nitrosamines over the 45 defined categories.

consistent, better informed and more scientifically rigorous. Many of the features of the prototype software tool presented and discussed in this paper will be available in Acrostic, a forthcoming read-across tool by Lhasa Limited. Acrostic will support use cases such as *N*-nitrosamine read-across to enable users to consider all important data, predictions and knowledge when evaluating potential analogues.

CRediT authorship contribution statement

Steven Kane: Writing – review & editing, Investigation, Data curation. **Dan Newman:** Writing – review & editing, Visualization, Software. **David J. Ponting:** Writing – review & editing, Investigation. **Edward Rosser:** Writing – review & editing, Software. **Robert Thomas:** Writing – review & editing, Methodology, Formal analysis. **Jonathan D. Vessey:** Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Conceptualization. **Samuel J. Webb:** Writing – review & editing, Software. **William H.J. Wood:** Writing – review & editing, Methodology.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data used in this article are available from Lhasa Limited.

Acknowledgements

We are grateful to Dr. Thierry Hanser and members of Lhasa's Molecular Informatics Team for the platform on which the prototype application was written, as well as numerous helpful discussions and support in use of the software. We are also grateful to Dr. Adrian Fowkes and Ally Reeves for their inputs in the preparation of the manuscript. In addition we would like to acknowledge the contribution of Dr. Fernanda Waechter for help with the validation of the approach described in this paper as well for reviewing and editing the manuscript.

The authors acknowledge the contribution of the many Lhasa Limited employees and contractors who have contributed to the collection of data and the development of knowledge and software which has been adapted for the purposes of this work.

References

[1] European Chemicals Agency, Read-Across Assessment Framework (RAAF), (n.d.). https://echa.europa.eu/documents/10162/13628/raaf_en.pdf/614e5d61-891d-4154-8a47-87efebd1851a (accessed August 23, 2022).

[2] S. Wu, K. Blackburn, J. Amburgey, J. Jaworska, T. Federle, A framework for using structural, reactivity, metabolic and physicochemical similarity to evaluate the suitability of analogs for SAR-based toxicological assessments, Regul. Toxicol. Pharm. 56 (2010) 67–81, https://doi.org/10.1016/J.YRTPH.2009.09.006.

[3] N. Ball, M.T.D. Cronin, J. Shen, K. Blackburn, E.D. Booth, M. Bouhifd, E. Donley, L. Egnash, C. Hastings, D.R. Juberg, A. Kleensang, N. Kleinstreuer, E.D. Kroese, A. C. Lee, T. Luechtefeld, A. Maertens, S. Marty, J.M. Naciff, J. Palmer, D. Pamies, M. Penman, A.N. Richarz, D.P. Russo, S.B. Stuard, G. Patlewicz, B. van Ravenzwaay, S. Wu, H. Zhu, T. Hartung, Toward Good Read-Across Practice (GRAP) guidance, ALTEX – Altern. Anim. Exper. 33 (2016) 149–166, https://doi.org/10.14573/ALTEX.1601251.

[4] S.E. Escher, A. Bitsch, Read-across methodology in toxicological risk assessment, Regul. Toxicol. (2021) 525–538, https://doi.org/10.1007/978-3-030-57499-4_132.

[5] I. Shah, J. Liu, R.S. Judson, R.S. Thomas, G. Patlewicz, Systematically evaluating read-across prediction and performance using a local validity approach characterized by chemical structure and bioactivity information, Regul. Toxicol. Pharm. 79 (2016) 12–24, https://doi.org/10.1016/J.YRTPH.2016.05.008.

[6] T.W. Schultz, P. Amcoff, E. Berggren, F. Gautier, M. Klaric, D.J. Knight, C. Mahony, M. Schwarz, A. White, M.T.D. Cronin, A strategy for structuring and reporting a read-across prediction of toxicity, Regul. Toxicol. Pharm. 72 (2015) 586–601, https://doi.org/10.1016/J.YRTPH.2015.05.016.

[7] C. Alexander-White, D. Bury, M. Cronin, M. Dent, E. Hack, N.J. Hewitt, G. Kenna, J. Naciff, G. Ouedraogo, A. Schepky, C. Mahony, C. Europe, A 10-step framework for use of read-across (RAX) in next generation risk assessment (NGRA) for cosmetics safety assessment, Regul. Toxicol. Pharm. 129 (2022) 105094, https://doi.org/10.1016/J.YRTPH.2021.105094.

[8] D. Ebbrell, M. Pator, M. Cronin, D8.2 Computational read-across workflows, (n.d.). https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5cacc7296&appId=PPGMS (accessed September 13, 2022).

[9] M. Masuda-Herrera, H.T. Rosen, A. Burild, T. Broschard, T. Griffin, T. Bell, J. Graham, J. Hillegass, P. Leavitt, B. Huta, T. Parris, U. Bruen, M. Cruz, J. Bercu, Harmonisation of read-across methodology for drug substance extractables and leachables (E&Ls), Regul. Toxicol. Pharm. (2023) 105494, https://doi.org/10.1016/J.YRTPH.2023.105494.

[10] G. Patlewicz, L.E. Lizarraga, D. Rua, D.G. Allen, A.B. Daniel, S.C. Fitzpatrick, N. Garcia-Reyero, J. Gordon, P. Hakkinen, A.S. Howard, A. Karmaus, J. Matheson, M. Mumtaz, A.N. Richarz, P. Ruiz, L. Scarano, T. Yamada, N. Kleinstreuer, Exploring current read-across applications and needs among selected U.S. Federal Agencies, Regulatory Toxicology and Pharmacology 106 (2019) 197–209. https://doi.org/10.1016/J.YRTPH.2019.05.011.

[11] G. Patlewicz, M.T.D. Cronin, G. Helman, J.C. Lambert, L.E. Lizarraga, I. Shah, Navigating through the minefield of read-across frameworks: A commentary perspective, Comput. Toxicol. 6 (2018) 39–54, https://doi.org/10.1016/J.COMTOX.2018.04.002.
[12] Health Canada, Use of analogues and read-across in risk assessment - Canada.ca, (n.d.). https://www.canada.ca/en/health-canada/services/chemical-substances/fact-sheets/analogues-read-across-risk-assessment.html (accessed September 13, 2022).
[13] C. Rovida, S.E. Escher, M. Herzler, S.H. Bennekou, H. Kamp, D.E. Kroese, L. Maslankiewicz, M.J. Moné, G. Patlewicz, N. Sipes, L. van Aerts, A. White, T. Yamada, B. van de Water, NAM-supported read-across: From case studies to regulatory guidance in safety assessment, ALTEX – Altern. Anim. Exper. 38 (2021) 140–150, https://doi.org/10.14573/ALTEX.2010062.
[14] S.E. Escher, H. Kamp, S.H. Bennekou, A. Bitsch, C. Fisher, R. Graepel, J. G. Hengstler, M. Herzler, D. Knight, M. Leist, U. Norinder, G. Ouédraogo, M. Pastor, S. Stuard, A. White, B. Zdrazil, B. van de Water, D. Kroese, Towards grouping concepts based on new approach methodologies in chemical hazard assessment: the read-across approach of the EU-ToxRisk project, Arch. Toxicol. 93 (2019) 3643–3667, https://doi.org/10.1007/S00204-019-02591-7/TABLES/5.
[15] C. Rovida, T. Barton-Maclaren, E. Benfenati, F. Caloni, P.C. Chandrasekera, C. Chesné, M.T.D. Cronin, J. de Knecht, D.R. Dietrich, S.E. Escher, S. Fitzpatrick, B. Flannery, M. Herzler, S.H. Bennekou, B. Hubesch, H. Kamp, J. Kisitu, N. Kleinstreuer, S. Kovarich, M. Leist, A. Maertens, K. Nugent, G. Pallocca, M. Pastor, G. Patlewicz, M. Pavan, O. Presgrave, L. Smirnova, M. Schwarz, T. Yamada, T. Hartung, Internationalization of read-across as a validated new approach method (NAM) for regulatory toxicology, ALTEX – Altern. Anim. Exper. 37 (2020) 579–606, https://doi.org/10.14573/ALTEX.1912181.
[16] G. Helman, I. Shah, G. Patlewicz, Extending the Generalised Read-Across approach (GenRA): A systematic analysis of the impact of physicochemical property information on read-across performance, Comput. Toxicol. 8 (2018) 34–50, https://doi.org/10.1016/J.COMTOX.2018.07.001.
[17] H. Zhu, M. Bouhifd, E. Donley, L. Egnash, N. Kleinstreuer, E.D. Kroese, Z. Liu, T. Luechtefeld, J. Palmer, D. Pamies, J. Shen, V. Strauss, S. Wu, T. Hartung, Supporting read-across using biological data, ALTEX – Altern. Anim. Exper. 33 (2016) 167–182, https://doi.org/10.14573/ALTEX.1601252.
[18] K. Blackburn, D. Bjerke, G. Daston, S. Felter, C. Mahony, J. Naciff, S. Robison, S. Wu, Case studies to test: A framework for using structural, reactivity, metabolic and physicochemical similarity to evaluate the suitability of analogs for SAR-based toxicological assessments, Regul. Toxicol. Pharm. 60 (2011) 120–135, https://doi.org/10.1016/J.YRTPH.2011.03.002.
[19] C.L. Mellor, R.L. Marchese Robinson, R. Benigni, D. Ebbrell, S.J. Enoch, J. W. Firman, J.C. Madden, G. Pawar, C. Yang, M.T.D. Cronin, Molecular fingerprint-derived similarity measures for toxicological read-across: Recommendations for optimal use, Regul. Toxicol. Pharm. 101 (2019) 121–134, https://doi.org/10.1016/J.YRTPH.2018.11.002.
[20] X. Zhang, S.P. Felter, A.M. Api, K. Joshi, D. Selechnik, A Cautionary tale for using read-across for cancer hazard classification: Case study of isoeugenol and methyl eugenol, Regul. Toxicol. Pharm. 136 (2022) 105280, https://doi.org/10.1016/J.YRTPH.2022.105280.
[21] A. Banerjee, M. Chatterjee, P. De, K. Roy, Quantitative predictions from chemical read-across and their confidence measures, Chemom. Intel. Lab. Syst. 227 (2022) 104613, https://doi.org/10.1016/J.CHEMOLAB.2022.104613.
[22] A. Worth, G. Tier, JRC Publications Repository - A Compendium of Case Studies that Helped to Shape the REACH Guidance on Chemical Categories and Read Across, JRC Publications Repository (2007). https://publications.jrc.ec.europa.eu/repository/handle/JRC37212 (accessed September 13, 2022).
[23] G. Patlewicz, N. Ball, E.D. Booth, E. Hulzebos, E. Zvinavashe, C. Hennes, Use of category approaches, read-across and (Q)SAR: General considerations, Regul. Toxicol. Pharm. 67 (2013) 1–12, https://doi.org/10.1016/J.YRTPH.2013.06.002.
[24] T. Yamada, T. Kawamura, S. Tsujii, M. Miura, H. Ohata, N. Katsutani, M. Matsumoto, A. Hirose, Formation and evaluation of mechanism-based chemical categories for regulatory read-across assessment of repeated-dose toxicity: A case of hemolytic anemia, Regul. Toxicol. Pharm. 136 (2022) 105275, https://doi.org/10.1016/J.YRTPH.2022.105275.
[25] K.P. Cross, D.J. Ponting, Developing structure-activity relationships for N-nitrosamine activity, Comput. Toxicol. 20 (2021) 100186, https://doi.org/10.1016/J.COMTOX.2021.100186.
[26] H.P. Gelbke, R. Ellis-Hutchings, H. Müllerschön, S. Murphy, M. Pemberton, Toxicological assessment of lower alkyl methacrylate esters by a category approach, Regul. Toxicol. Pharm. 92 (2018) 104–127, https://doi.org/10.1016/J.YRTPH.2017.11.013.
[27] G. Helman, I. Shah, A.J. Williams, J. Edwards, J. Dunne, G. Patlewicz, Generalized Read-Across (GenRA): A workflow implemented into the EPA CompTox Chemicals Dashboard, ALTEX 36 (2019) 462–465, https://doi.org/10.14573/ALTEX.1811292.
[28] G. Gini, A.M. Franchi, A. Manganaro, A. Golbamaki, E. Benfenati, ToxRead: A tool to assist in read across and its use to assess mutagenicity of chemicals, 25 (2014) 999–1011. https://doi.org/10.1080/1062936X.2014.976267.
[29] S. Kutsarova, T.W. Schultz, A. Chapkanov, D. Cherkezova, A. Mehmed, S. Stoeva, C. Kuseva, D. Yordanova, M. Georgiev, T. Petkov, O.G. Mekenyan, The QSAR Toolbox automated read-across workflow for predicting acute oral toxicity: II. Verification and Validation, Comput. Toxicol. 20 (2021) 100194, https://doi.org/10.1016/J.COMTOX.2021.100194.
[30] S. Kutsarova, A. Mehmed, D. Cherkezova, S. Stoeva, M. Georgiev, T. Petkov, A. Chapkanov, T.W. Schultz, O.G. Mekenyan, Automated read-across workflow for predicting acute oral toxicity: I. The Decision Scheme in the QSAR Toolbox, Regul. Toxicol. Pharmacol. 125 (2021) 105015, https://doi.org/10.1016/J.YRTPH.2021.105015.
[31] A.Y. Caballero Alfonso, C. Chayawan, D. Gadaleta, A. Roncaglioni, E. Benfenati, A KNIME workflow to assist the analogue identification for read-across, applied to aromatase activity, Molecules 28 (2023) 1832, https://doi.org/10.3390/MOLECULES28041832/S1.
[32] H. Moustakas, M.S. Date, M. Kumar, T.W. Schultz, D.C. Liebler, T.M. Penning, D. T. Salvito, A.M. Api, An end point-specific framework for read-across analog selection for human health effects, Chem. Res. Toxicol. 35 (2022) 2324–2334, https://doi.org/10.1021/ACS.CHEMRESTOX.2C00286/ASSET/IMAGES/LARGE/TX2C00286_0007.JPEG.
[33] C. Lester, E.L. Byrd, M. Shobair, G. Yan, Quantifying analogue suitability for SAR-based read-across toxicological assessment, Chem Res Toxicol 36 (2023) 230–242, https://doi.org/10.1021/ACS.CHEMRESTOX.2C00311/SUPPL_FILE/TX2C00311_SI_002.XLS.
[34] QSAR Flex, (n.d.). https://www.multicase.com/qsar-flex (accessed November 26, 2022).
[35] Instem Unveils New Additions to its Leading Computational Toxicology Software Suite, (n.d.). https://www.instem.com/news/articles/0211-new-additions-to-computational-toxicology-software-suite.php (accessed November 12, 2022).
[36] EMA reviewing medicines containing valsartan from Zhejiang Huahai following detection of an impurity: some valsartan medicines being recalled across the EU | European Medicines Agency, (n.d.). https://www.ema.europa.eu/en/news/ema-reviewing-medicines-containing-valsartan-zhejiang-huahai-following-detection-impurity-some (accessed November 26, 2022).
[37] Pfizer Expands Voluntary Nationwide Recall to include All Lots of CHANTIX® (Varenicline) Tablets Due to N-Nitroso Varenicline Content | FDA, (n.d.). https://www.fda.gov/safety/recalls-market-withdrawals-safety-alerts/pfizer-expands-voluntary-nationwide-recall-include-all-lots-chantixr-varenicline-tablets-due-n#recall-announcement (accessed November 26, 2022).
[38] J.E. Doe, A.R. Boobis, V. Dellarco, P.A. Fenner-Crisp, A. Moretto, T.P. Pastoor, R. S. Schoeny, J.G. Seed, D.C. Wolf, Chemical carcinogenicity revisited 2: Current knowledge of carcinogenesis shows that categorization as a carcinogen or non-carcinogen is not scientifically credible, Regul. Toxicol. Pharm. 103 (2019) 124–129, https://doi.org/10.1016/J.YRTPH.2019.01.024.
[39] European Medicines Agency, Assessment report. Procedure under Article 5(3) of Regulation EC (No) 726/2004. Nitrosamine impurities in human medicinal products, (n.d.). https://www.ema.europa.eu/en/documents/referral/nitrosamines-emea-h-a53-1490-assessment-report_en.pdf (accessed January 10, 2024).
[40] Lhasa Limited, Lhasa Carcinogenicity Database, (n.d.). https://www.lhasalimited.org/products/lhasa-carcinogenicity-database.htm (accessed August 24, 2022).
[41] G.E. Johnson, K. Dobo, B. Gollapudi, J. Harvey, J. Kenny, M. Kenyon, A. Lynch, S. Minocherhomji, J. Nicolette, V. Thybaud, R. Wheeldon, A. Zeller, Permitted daily exposure limits for noteworthy N-nitrosamines, Environ. Mol. Mutagen 62 (2021) 293–305, https://doi.org/10.1002/EM.22446.
[42] A. Thresher, R. Foster, D.J. Ponting, S.A. Stalford, R.E. Tennant, R. Thomas, Are all nitrosamines concerning? A review of mutagenicity and carcinogenicity data, Regul. Toxicol. Pharm. 116 (2020) 104749, https://doi.org/10.1016/J.YRTPH.2020.104749.
[43] C.E.M. dos Santos, D.J. Dorta, D.P. de Oliveira, Setting limits for N-nitrosamines in drugs: A defined approach based on read-across and structure-activity relationship for N-nitrosopiperazine impurities, Regul. Toxicol. Pharm. 136 (2022) 105288, https://doi.org/10.1016/J.YRTPH.2022.105288.
[44] European Medicines Agency, European Medicines Regulatory Network approach for the implementation of the CHMP Opinion pursuant to Article 5(3) of Regulation (EC) No 726/2004 for nitrosamine impurities in human medicines, (n.d.). https://www.ema.europa.eu/en/documents/referral/european-medicines-regulatory-network-approach-implementation-chmp-opinion-pursuant-article-53-regulation-ec-no-726-2004-nitrosamine-impurities-human-medicines_en.pdf (accessed January 10, 2024).
[45] European Medicines Agency, Questions and answers for marketing authorisation holders/applicants on the CHMP Opinion for the Article 5(3) of Regulation (EC) No 726/2004 referral on nitrosamine impurities in human medicinal products, (n.d.). https://www.ema.europa.eu/en/documents/referral/nitrosamines-emea-h-a53-1490-questions-answers-marketing-authorisation-holders/applicants-chmp-opinion-article-53-regulation-ec-no-726/2004-referral-nitrosamine-impurities-human-medicinal-products_en.pdf (accessed August 30, 2023).
[46] Carcinogenicity Assessment | Lhasa Limited, (n.d.). https://www.lhasalimited.org/solutions/carcinogenicity-assessment/ (accessed November 13, 2023).
[47] Nitrosamine Impurity Risk Assessment | Lhasa Limited, (n.d.). https://www.lhasalimited.org/solutions/nitrosamine-impurity-risk-assessment/ (accessed November 13, 2023).
[48] A. Thresher, J.P. Gosling, R. Williams, Generation of TD50 values for carcinogenicity study data, Toxicol. Res. (Camb.) 8 (2019) 696–703, https://doi.org/10.1039/C9TX00118B.
[49] D.J. Ponting, E. Murray, A. Long, Quantifying confidence in the reporting of metabolic biotransformations, Drug Discov. Today 22 (2017) 970–975, https://doi.org/10.1016/J.DRUDIS.2017.01.001.
[50] T. Ball, C.G. Barber, A. Cayley, M.L. Chilton, R. Foster, A. Fowkes, C. Heghes, E. Hill, N. Hill, S. Kane, D.S. Macmillan, A. Myden, D. Newman, A. Polit, S. A. Stalford, J.D. Vessey, Beyond adverse outcome pathways: making toxicity predictions from event networks, SAR models, data and knowledge, Toxicol. Res. (Camb.) 10 (2021) 102–122, https://doi.org/10.1093/TOXRES/TFAA099.

[51] R. Thomas, R.E. Tennant, A.A.F. Oliveira, D.J. Ponting, What makes a potent nitrosamine? Statistical validation of expert-derived structure activity relationships, Chem. Res. Toxicol. 35 (2022) 1997–2013.
[52] D.J. Ponting, K.L. Dobo, M.O. Kenyon, A.S. Kalgutkar, On the Carcinogenic Potential of Novel N-Nitrosamine Impurities Derived from Active Pharmaceutical Ingredients, J. Med. Chem. in press (n.d.).
[53] D. Ponting, SAR Development and MOA analysis to help set AIs, in: FDA/HESI Research Roadmap Planning on Hazard and Risk Assessment of Nitrosamine Impurities in Drugs, Washington, DC, 2023.
[54] C.A. Marchant, K.A. Briggs, A. Long, In silico tools for sharing data and knowledge on toxicity and metabolism: Derek for windows, meteor, and vitic, 18 (2008) 177–187. https://doi.org/10.1080/15376510701857320.
[55] Metabolite Identification and Analysis | Lhasa Limited, (n.d.). https://www.lhasalimited.org/solutions/metabolite-identification-and-analysis/ (accessed November 13, 2023).
[56] D. Mendez, A. Gaulton, A.P. Bento, J. Chambers, M. de Veij, E. Félix, M.P. Magariños, J.F. Mosquera, P. Mutowo, M. Nowotka, M. Gordillo-Marañón, F. Hunter, L. Junco, G. Mugumbate, M. Rodriguez-Lopez, F. Atkinson, N. Bosc, C.J. Radoux, A. Segura-Cabrera, A. Hersey, A.R. Leach, ChEMBL: towards direct deposition of bioassay data, Nucleic Acids Res. 47 (2019) D930–D940. https://doi.org/10.1093/NAR/GKY1075.
[57] F. Pedregosa, V. Michel, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, J. Vanderplas, D. Cournapeau, F. Pedregosa, G. Varoquaux, A. Gramfort, B. Thirion, O. Grisel, V. Dubourg, A. Passos, M. Brucher, M. Perrot, É. Duchesnay, Scikit-learn: Machine Learning in Python, J. Mach. Learn. Res. 12 (2011) 2825–2830, https://doi.org/10.5555/1953048.
[58] G. Landrum, RDKit, Https://Www.Rdkit.Org/ (2010). https://www.rdkit.org/ (accessed November 30, 2022).
[59] H. Papadopoulos, V. Vovk, A. Gammerman, Regression conformal prediction with nearest neighbours, J. Artif. Intell. Res. 40 (2011) 815–840, https://doi.org/10.1613/JAIR.3198.
[60] A. Sartini, Prediction of water solubility of organic compounds using Lhasa in-house descriptors, in: 22nd European Symposium on Quantitative Structure-Activity Relationships (EuroQSAR), Thessaloniki, Greece, 2018.
[61] J. Plante, S. Werner, JPlogP: An improved logP predictor trained using predicted data, J. Cheminform. 10 (2018) 1–10, https://doi.org/10.1186/S13321-018-0316-5/TABLES/4.
[62] T. Hanser, Personal communication, Personal Communication (n.d.).
[63] A.N. Cayley, R.S. Foster, E. Hill, S. Kane, G. Kocks, A. Myden, D. Newman, S.A. Stalford, J.D. Vessey, R. Zarei, A. Anax, F. de Oliveira, Development of a network of carcinogenicity adverse outcome pathways and its employment as an evidence framework for safety assessment, ALTEX - Alternatives to Animal Experimentation (2022). https://doi.org/10.14573/ALTEX.2201311.
[64] Lhasa Limited, Kaptis, (n.d.). https://www.lhasalimited.org/products/kaptis.htm (accessed August 23, 2022).
[65] S.M. Cohen, A.R. Boobis, V.L. Dellarco, J.E. Doe, P.A. Fenner-Crisp, A. Moretto, T. P. Pastoor, R.S. Schoeny, J.G. Seed, D.C. Wolf, Chemical carcinogenicity revisited 3: Risk assessment of carcinogenic potential based on the current state of knowledge of carcinogenesis in humans, Regul. Toxicol. Pharm. 103 (2019) 100–105, https://doi.org/10.1016/J.YRTPH.2019.01.017.
[66] A.R. Boobis, S.M. Cohen, V.L. Dellarco, J.E. Doe, P.A. Fenner-Crisp, A. Moretto, T. P. Pastoor, R.S. Schoeny, J.G. Seed, D.C. Wolf, Classification schemes for carcinogenicity based on hazard-identification have become outmoded and serve neither science nor society, Regul. Toxicol. Pharm. 82 (2016) 158–166, https://doi.org/10.1016/J.YRTPH.2016.10.014.
[67] S.P. Felter, X. Zhang, C. Thompson, Butylated hydroxyanisole: Carcinogenic food additive to be avoided or harmless antioxidant important to protect food supply? Regul. Toxicol. Pharm. 121 (2021) 104887 https://doi.org/10.1016/J.YRTPH.2021.104887.
[68] European Medicines Agency, Appendix 2 to Questions and answers for marketing authorisation holders/applicants on the CHMP Opinion for the Article 5(3) of Regulation (EC) No 726/2004 referral on nitrosamine impurities in human medicinal products, (n.d.). https://www.ema.europa.eu/en/documents/other/appendix-2-questions-and-answers-marketing-authorisation-holders-applicants-chmp-opinion-article-53-regulation-ec-no-726-2004-referral-nitrosamine-impurities-human-medicinal-products_en.pdf (accessed January 9, 2024).
[69] K.L. Dobo, M.O. Kenyon, O. Dirat, M. Engel, A. Fleetwood, M. Martin, S. Mattano, A. Musso, J.C. McWilliams, A. Papanikolaou, P. Parris, J. Whritenour, S. Yu, A. S. Kalgutkar, Practical and science-based strategy for establishing acceptable intakes for drug product N-nitrosamine impurities, Chem. Res. Toxicol. 35 (2022) 475–489, https://doi.org/10.1021/ACS.CHEMRESTOX.1C00369/ASSET/IMAGES/LARGE/TX1C00369_0004.JPEG.
[70] C.A. Marchant, E.M. Rosser, J.D. Vessey, A k-nearest neighbours approach using metabolism-related fingerprints to improve in silico metabolite ranking, Mol. Inform. 36 (2017) 1600105, https://doi.org/10.1002/MINF.201600105.
[71] US Food and Drug Administration, FDA Updates and Press Announcements on Nitrosamine in Varenicline (Chantix) | FDA, (n.d.). https://www.fda.gov/drugs/drug-safety-and-availability/fda-updates-and-press-announcements-nitrosamine-varenicline-chantix#632a7e55b9e11 (accessed September 21, 2022).
[72] A. Long, S. Webb, Beyond Simple Chemical Structure Similarity Assessment (2) Automation of Biotransformation Profile Similarity Calculation and Display, in: 5th EU-ToxRisk General Assembly Meeting, Egmond aan Zee, NL, 2019.
[73] S. Riniker, G.A. Landrum, Similarity maps - A visualization strategy for molecular fingerprints and machine-learning methods, J. Cheminform. 5 (2013) 1–7, https://doi.org/10.1186/1758-2946-5-43/FIGURES/5.
[74] OECD, Test No. 471: Bacterial Reverse Mutation Test, OECD, 2020. https://doi.org/10.1787/9789264071247-en.
[75] R.V. Williams, D.M. DeMarini, L.F. Stankowski, P.A. Escobar, E. Zeiger, J. Howe, R. Elespuru, K.P. Cross, Are all bacterial strains required by OECD mutagenicity test guideline TG471 needed? Mutat. Res./Genet. Toxicol. Environ. Mutagen. 848 (2019) 503081 https://doi.org/10.1016/J.MRGENTOX.2019.503081.
[76] A.W. Andrews, W. Lijinsky, The mutagenicity of 45 nitrosamines in Salmonella typhimurium, Teratog. Carcinog. Mutagen. 1 (1981) 295–303, https://doi.org/10.1002/TCM.1770010306.
[77] Nitrosamine impurities | European Medicines Agency, (n.d.). https://www.ema.europa.eu/en/human-regulatory/post-authorisation/referral-procedures/nitrosamine-impurities (accessed October 23, 2023).
[78] Nitrosamine impurities in medications: Guidance - Canada.ca, (n.d.). https://www.canada.ca/en/health-canada/services/drugs-health-products/compliance-enforcement/information-health-product/drugs/nitrosamine-impurities/medications-guidance.html (accessed October 23, 2023).
[79] European Medicines Agency, Committee for Medicinal Products for Human Use (CHMP) Assessment report, 2019. https://www.ema.europa.eu/en/documents/variation-report/angiotensin-ii-receptor-antagonists-sartans-article-31-referral-chmp-assessment-report_en.pdf (accessed September 22, 2022).
[80] Lhasa Limited, Vitic, Lhasa Limited (n.d.). https://www.lhasalimited.org/products/vitic.htm (accessed September 16, 2022).
[81] S. Li, L. Dong, K. Tang, Z. Lan, R. Liu, Y. Wang, R. Wang, H. Lin, Simultaneous and trace level quantification of two potential genotoxic impurities in valsartan drug substance using UPLC-MS/MS, J. Pharm. Biomed. Anal. 212 (2022) 114630, https://doi.org/10.1016/J.JPBA.2022.114630.
[82] A. Amberg, J.S. Harvey, A. Czich, H.P. Spirkl, S. Robinson, A. White, D.P. Elder, Do carboxylic/sulfonic acid halides really present a mutagenic and carcinogenic risk as impurities in final drug products? Org. Process. Res. Dev. 19 (2015) 1495–1506, https://doi.org/10.1021/ACS.OPRD.5B00106/SUPPL_FILE/OP5B00106_SI_001.PDF.
[83] A. Trejo-Martin, J.P. Bercu, A. Thresher, R.E. Tennant, R.F. Thomas, K. Cross, A. Czich, K. Waese, J.J. Nicolette, J. Murray, P. Sonders, A. Kondratiuk, J. R. Cheung, D. Thomas, A. Lynch, J. Harvey, S. Glowienke, L. Custer, P.A. Escobar, Use of the bacterial reverse mutation assay to predict carcinogenicity of N-nitrosamines, Regul. Toxicol. Pharm. 135 (2022) 105247, https://doi.org/10.1016/J.YRTPH.2022.105247.
[84] S. Glowienke, U. Onken, A. Elhajouji, S. Muthusamy, R. Sangana, H.J. Martus, T. Bedman, A. Hartmann, Genotoxicity evaluation of a valsartan-related complex N-nitroso-impurity, Regul. Toxicol. Pharm. 134 (2022) 105245, https://doi.org/10.1016/J.YRTPH.2022.105245.