



Guide to Mutagenicity Classification Model version 2.1.13

Table of Contents

1. Model explanation.....	2
1.1 Introduction.....	2
1.2 Model details.....	2
1.3 Applicability Domain.....	3
1.4 Structural Alerts for mutagen compounds.....	5
1.5 Structural Alerts for suspect mutagen compounds.....	5
1.6 Model statistics.....	6
2. Model usage.....	6
2.1 Input.....	6
2.2 Output.....	7
3. Differences from previous versions.....	8
3.1 Differences with CAESAR model.....	8
3.2 VEGA model history.....	8
3.2.1 Version 2.1.6.....	8
3.2.2 Version 2.1.7.....	8
3.2.3 Version 2.1.8.....	8
3.2.4 Version 2.1.9.....	9
3.2.5 Version 2.1.10.....	9
3.2.6 Version 2.1.12.....	9
3.2.7 Version 2.1.13.....	9

1. Model explanation

1.1 Introduction

The model provides a qualitative prediction of mutagenicity on *Salmonella typhimurium* (Ames test). It is implemented inside the VEGA online platform, accessible at: <http://www.vega-qsar.eu/>. The model extends the original CAESAR model, freely available at: <http://www.caesar-project.eu/software/>

1.2 Model details

An integrated model was arranged cascading two models: Model A, a trained Support Vector Machine (SVM) classifier, and an additional Model B for false negatives (FNs) removal based on Structural Alerts (SAs) matching. Model B works with two sets of SAs: the first one is related to mutagenicity activity, if no matching is found the second set is checked, and if some matches are found, the prediction is "suspect mutagen". Full reference and details of the used formulas can be found in:

Ferrari T., Gini G. An open source multistep model to predict mutagenicity from statistical analysis and relevant structural alerts. *Chemistry Central Journal* (2010), 4 (Suppl 1):S2

Structural Alerts have been taken from the Benigni/Bossa rulebase for mutagenicity and carcinogenicity (implemented as a module of Toxtree software). Full reference and details of the used SAs can be found in:

Benigni R., Bossa C., Jeliaskova N.G., Netzeva T.I., Worth A.P. The Benigni/Bossa rulebase for mutagenicity and carcinogenicity - a module of toxtree. Technical Report EUR 23241 EN, European Commission - Joint Research Centre 2008.

The descriptors used by the SVM are the following:

- SsCH3_acnt: Count of all (– CH3) groups in molecule
- SdCH2_acnt: Count of all (= CH2) groups in molecule
- SssCH2_acnt: Count of all (– CH2 –) groups in molecule
- SdsCH_acnt: Count of all (= CH –) groups in molecule
- SaaCH_acnt: Count of all (CH) groups in molecule
- SsssCH_acnt: Count of all (> CH –) groups in molecule
- SdssC_acnt: Count of all (= C <) groups in molecule
- SaasC_acnt: Count of all (CH) groups in molecule
- SaaaC_acnt: Count of all (CH) groups in molecule
- SssssC_acnt: Count of all (> C <) groups in molecule
- SsNH2_acnt: Count of all (– NH2) groups in molecule
- StN_acnt: Count of all (#N) groups in molecule
- SdsN_acnt: Count of all (= N –) groups in molecule
- SaaN_acnt: Count of all (N) groups in molecule

- SsssN_acnt: Count of all (> N –) groups in molecule
- SdaaN_acnt: Count of all (N) groups in molecule
- SsOH_acnt: Count of all (– OH) groups in molecule
- SdO_acnt: Count of all (= O) groups in molecule
- SssO_acnt: Count of all (– O –) groups in molecule
- SaaO_acnt: Count of all (O) groups in molecule
- SHCHnX_Acnt: Count of all CH or CH₂ groups with a -F or -Cl also bonded to the carbon
- Gmin: Smallest atom E-State value in molecule
- idwbar: Bonchev-Trinajsti mean information content
- ALogP: Ghose-Crippen octanol water coefficient
- nrings: Number of rings (cyclomatic number) in a molecular graph

The descriptors were calculated, in the original model, by means of MDL and dragonX software and are now entirely calculated by an in-house software module in which they are implemented as described in: R. Todeschini and V. Consonni, *Molecular Descriptors for Chemoinformatics*, Wiley-VCH, 2009.

1.3 Applicability Domain

The applicability domain of predictions is assessed using an Applicability Domain Index (ADI) that has values from 0 (worst case) to 1 (best case). The ADI is calculated by grouping several other indices, each one taking into account a particular issue of the applicability domain. Most of the indices are based on the calculation of the most similar compounds found in the training and test set of the model, calculated by a similarity index that consider molecule's fingerprint and structural aspects (count of atoms, rings and relevant fragments).

For each index, including the final ADI, three intervals for its values are defined, such that the first interval corresponds to a positive evaluation, the second one corresponds to a suspicious evaluation and the last one corresponds to a negative evaluation.

Following, all applicability domain components are reported along with their explanation and the intervals used. Note that for purpose of evaluating accuracy and concordance indices, prediction of "suspect mutagen" are considered as "mutagen".

- **Similar molecules with known experimental value.** This index takes into account how similar are the first three most similar compounds found. Values near 1 mean that the predicted compound is well represented in the dataset used to build the model, otherwise the prediction could be an extrapolation. Defined intervals are:

1 >= index > 0.85	strongly similar compounds with known experimental value in the training set have been found
0.85 >= index > 0.7	only moderately similar compounds with known experimental value in the training set have been found
index <= 0.7	no similar compounds with known experimental value in the training set have been found

- **Accuracy of prediction for similar molecules.** This index takes into account the classification accuracy in prediction for the three most similar compounds found. Values near 1 mean that the predicted compounds falls in an area of the model's space where the model gives reliable predictions (no misclassifications), otherwise the lower is the value, the worse the model behaves. Defined intervals are:

$1 \geq \text{index} > 0.9$	accuracy of prediction for similar molecules found in the training set is good
$0.9 \geq \text{index} > 0.5$	accuracy of prediction for similar molecules found in the training set is not optimal
$\text{index} \leq 0.5$	accuracy of prediction for similar molecules found in the training set is not adequate

- **Concordance for similar molecules** . This index takes into account the difference between the predicted value and the experimental values of the three most similar compounds. Values near 0 mean that the prediction made disagrees with the values found in the model's space, thus the prediction could be unreliable. Defined intervals are:

$1 \geq \text{index} > 0.9$	similar molecules found in the training set have experimental values that agree with the predicted value
$0.9 \geq \text{index} > 0.5$	some similar molecules found in the training set have experimental values that disagree with the predicted value
$\text{index} \leq 0.5$	similar molecules found in the training set have experimental values that disagree with the predicted value

- **Atom Centered Fragments similarity check.** This index takes into account the presence of one or more fragments that aren't found in the training set, or that are rare fragments. First order atom centered fragments from all molecules in the training set are calculated, then compared with the first order atom centered fragments from the predicted compound; then the index is calculated as following: a first index RARE takes into account rare fragments (those who occur less than three times in the training set), having value of 1 if no such fragments are found, 0.85 if up to 2 fragments are found, 0.7 if more than 2 fragments are found; a second index NOTFOUND takes into account not found fragments, having value of 1 if no such fragments are found, 0.6 if a fragments is found, 0.4 if more than 1 fragment is found. Then, the final index is given as the product $\text{RARE} * \text{NOTFOUND}$. Defined intervals are:

$\text{index} = 1$	all atom centered fragment of the compound have been found in the compounds of the training set
$1 > \text{index} \geq 0.7$	some atom centered fragment of the compound have not been found in the compounds of the training set or are rare fragments
$\text{index} < 0.7$	a prominent number of atom centered fragments of the compound have not been found in the compounds of the training set or are rare fragments

- **Model descriptors range check.** This index checks if the descriptors calculated for the predicted compound are inside the range of descriptors of the training and test set. The index has value 1 if all descriptors are inside the range, 0 if at least one descriptor is out of the range. Defined intervals are:

index = True	descriptors for this compound have values inside the descriptor range of the compounds of the training set
index = False	descriptors for this compound have values outside the descriptor range of the compounds of the training set

- **Global AD Index.** The final global index takes into account all the previous indices, in order to give a general global assessment on the applicability domain for the predicted compound. Defined intervals are:

$1 \geq \text{index} \geq 0.9$	predicted substance is into the Applicability Domain of the model
$0.9 > \text{index} \geq 0.7$	predicted substance could be out of the Applicability Domain of the model
$\text{index} < 0.7$	predicted substance is out of the the Applicability Domain of the model

1.4 Structural Alerts for mutagen compounds

The model goes through a first step in which a set of 12 SAs related to mutagenicity is checked. If one (or more) fragments are found, compound is predicted as "mutagen". The SAs are the following (SA numbers refer to the original Benigni/Bossa study):

- SA 1: Acyl halides
- SA 6: Propiolactones or propiosultones
- SA 12: Quinones
- SA 13: Hydrazine
- SA 14: Aliphatic azo and azoxy
- SA 16: alkyl carbamate and thiocarbamate
- SA 18: Polycyclic Aromatic Hydrocarbons
- SA 21: alkyl and aryl N-nitroso groups
- SA 22: Azide and triazene groups
- SA 25: Aromatic nitroso group
- SA 28bis: Aromatic mono- and dialkylamine
- SA 29: Aromatic diazo

1.5 Structural Alerts for suspect mutagen compounds

After the first step, if no mutagen fragment has been found, the model checks a second set of 4 SAs related to suspect mutagenicity. In fact, these SAs have a moderate rate of false positive in the training set of the model. If one (or more) fragments are found, compound is predicted as "suspect mutagen". The SAs are the following (SA numbers refer to the original Benigni/Bossa study):

- SA 7: Epoxides and aziridines
- SA 8: Aliphatic halogens
- SA 19: Heterocyclic Polycyclic Aromatic Hydrocarbons
- SA 27: Nitro-aromatic

1.6 Model statistics

Following, statistics obtained applying the model to its original dataset:

- Training set: n = 3253; Accuracy = 0.92; Specificity = 0.86; Sensitivity = 0.97
- Test set: n = 798; Accuracy = 0.83; Specificity = 0.74; Sensitivity = 0.90

From these statistics, compounds predicted as Suspicious Mutagen have been omitted. Structural alerts for suspect mutagenicity have been found as follows:

- In the training set: n = 114 (18 mutagen compounds, 96 non-mutagen compounds)
- In the test set: n = 39 (19 mutagen compounds, 20 non-mutagen compounds)

2. Model usage

2.1 Input

The model accepts as input two molecule formats: SDF (multiple MOL file) and SMILES. All molecules found as input are preprocessed before the calculation of molecular descriptors, in order to obtain a standardized representation of compound. For this reason, some cautions should be taken.

- **Hydrogen atoms.** In SDF files, hydrogen atoms should be explicit. As some times SDF file store only skeleton atoms, and hydrogen atoms are implicit, during the processing of the molecule the system tries to add implicit hydrogens on the basis of the known standard valence of each atom (for example, if a carbon atoms has three single bonds, an hydrogen atom will be added such to reach a valence of four). In SMILES molecules, the default notation uses implicit hydrogen. Anyway please note that in some cases it is necessary to explicitly report an hydrogen; this happens when the conformation is not unambiguous. For example, when a nitrogen atom is into an aromatic ring with a notation like "cnc" it is not clear whether it corresponds to C-N=C or to C-[NH]-C, thus if the situation is the latter, it should be explicitly reported as "c[nH]c".

- **Aromaticity.** The system calculates aromaticity using the basic Hueckel rule. Note that each software for drawing and storing of molecules can use different approaches to aromaticity (for instance, commonly the user can choose between the basic Hueckel rule and a loose approach that lead to considering aromatic a greater number of rings). As in the input files aromaticity can be set explicitly (for instance, in SMILES format by using lowercase letters), during the processing of the molecule the system removes aromaticity from rings that don't satisfy the Hueckel rule. Please note that when aromaticity is removed from a ring, it is not always possible to rebuild the original structure in Kekule form (i.e. with an alternation of single and double bonds, like in the SMILES for benzene, C=1C=CC=CC1), in this case all bonds are set to single. Furthermore, please note that aromaticity detection is a really relevant issue, some molecular descriptors can have significantly different values whether a ring is perceived as aromatic or not. For this reason it is strongly recommended:

- Always use explicit hydrogens in SDF file.
- Avoid explicit aromaticity notation in original files; in this way, the perception of aromaticity is left to the preprocessing step and there is no chance of mistakes due to the transformation of rings that were set to aromatic in the original format but not recognized as aromatic in VEGA.

Note that when some modification of the molecule are performed during the preprocessing (e.g. adding of lacking hydrogens, correction of aromaticity), a warning is given in the remark field of the results.

2.2 Output

Results given as text file consist of a plain-text tabbed file (easily importable and processable by any spreadsheet software) containing in each row all the information about the prediction of a molecule. Note that if some problems were encountered while processing the molecule structure, some warnings are reported in the last field (Remarks).

Results given as PDF file consists of a document containing all the information about the prediction. For each molecule, results are organized in sections with the following order:

1 – Prediction summary

Here is reported a depiction of the compound and the final assessment of the prediction (i.e. the prediction made together with the analysis of the applicability domain). Note that if some problems were encountered while processing the molecule structure, some warnings are reported in the last field (Remarks).

A graphical representation of the evaluation of the prediction and of its reliability is also provided, using the following elements:



Compound is classified as non-mutagen



Compound is classified as suspect mutagen



Compound is classified as mutagen



Prediction has low reliability (compound out of the AD)



Prediction has moderate reliability (compound could be out of the AD)



Prediction has high reliability (compound into the AD)

3.1 – Applicability Domain: Similar compounds, with predicted and experimental values

Here it is reported the list of the six most similar compounds found in the training and test set of the model, along with their depiction and relevant information (mainly experimental value and predicted value).

3.2 – *Applicability Domain: Measured Applicability Domain scores*

Here it is reported the list of all Applicability Domain scores, starting with the global Applicability Domain Index (ADI). Note that the final assessment on prediction reliability is given on the basis of the value of the ADI. For each index, it is reported its value and a brief explanation of the meaning of that value.

4.1 – *Reasoning: Relevant chemical fragments and moieties*

If some rare and/or missing Atom Centered Fragments are found, they are reported here with a depiction of each fragment.

If some relevant fragments are found (see section 1.4 and 1.5 of this guide), they are reported here (one for each page) with a brief explanation of their meaning and the list of the three most similar compounds that contain the same fragment. Note that if no relevant fragments are found, this section is not shown.

3. Differences from previous versions

3.1 Differences with CAESAR model

The VEGA model has several differences that can lead to prediction that can be slightly different from the ones produced by the CAESAR model. Mainly, descriptors are no longer calculated by means of third party software thus, even if the algorithm definition is the same, their values could be different for some molecules. Please note that also the algorithm for the calculation of similarity is slightly different from the one implemented in CAESAR.

3.2 VEGA model history

3.2.1 Version 2.1.6

First official release published in the VEGA platform.

3.2.2 Version 2.1.7

Some minor code updates, mainly due to changes in the core. There are NO changes in prediction values and AD assessment.

3.2.3 Version 2.1.8

This version is updated with the new calculation core (1.0.26) where similarity algorithm is slightly changed. The new version considers halogen atoms are really similar, especially Chlorine and Bromine atoms are considered almost the same. The main difference with previous algorithm can be thus seen just for halogenated compounds.

A more precise check for similarity has been introduced for the extraction of experimental values, in order to avoid mismatches (as the similarity index is based on fingerprints, there are some rare cases in which a value equal to 1 does not point to an exactly isomorph compound).

Some minor bugs in the procedure for reading molecule structures have been fixed; some compounds, previously not loaded, could now be correctly processed.

There are NO changes in prediction values, but as similarity is changed some small differences in AD assessment can be found.

3.2.4 Version 2.1.9

This version is updated with the new calculation core (1.0.27), that generates a graphically renewed PDF report. In this version, the propositions for prediction and assessment are changed, but there are NO changes in their values.

3.2.5 Version 2.1.10

This version is updated with the new calculation core (1.0.30).

In this version, a major bug was fixed: molecule structures were modified while calculating some descriptors, thus when this model was executed together with other models (in VegaNIC), the results of the models executed after this one could be affected by some errors.

3.2.6 Version 2.1.12

This version is updated with the new calculation core (1.1.1) based on a new release of the CDK libraries (1.4.9). These updates can influence the calculation, so there could be some changes in the predictions produced.

The Structural Alerts implementation have been revised and corrected, there could be changes in the SAs found and thus also in the prediction of some compounds.

The new calculation core implements a new version of the algorithm used for calculating the similarity index. This means that the list of similar molecules given as part of the applicability domain evaluation will often be different from the ones produced by older releases of the model. Furthermore, the applicability domain index (ADI) itself and the final assessment could often be different.

Model statistics in the current guide have been updated with the new values.

3.2.7 Version 2.1.13

This version is updated with the new calculation core (1.2.0). This update can influence some calculation, in particular similarity evaluation, so there could be some changes in the applicability domain values produced.

Furthermore, some predictions can be different from previous versions as in the new calculation core Benigni/Bossa alerts have been completely revised (using ToxTree 2.6.13 as reference for their implementation)