

Advanced analytics platform

The power of enhanced analytics and machine learning

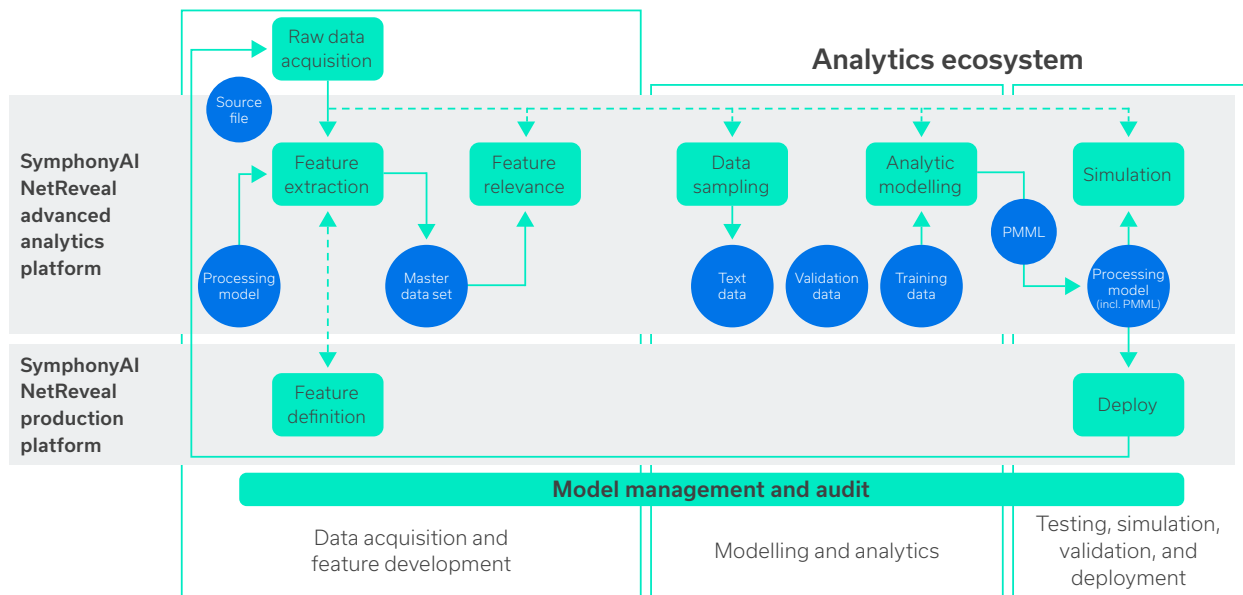
SymphonyAI NetReveal advanced analytics platform

The advanced analytics platform (AAP) is SymphonyAI NetReveal's specialist financial crime analytics and machine learning engine. It is driven by data science workflows essential to deliver improved detection performance in terms of increased effectiveness and efficiency. AAP is easy to use, does not require users to have deep data science domain skills, and requires no complex code or scripts to be written.

AAP adds enhanced analytics and machine learning across the SymphonyAI NetReveal product suite and has two main interactions with the core platform:

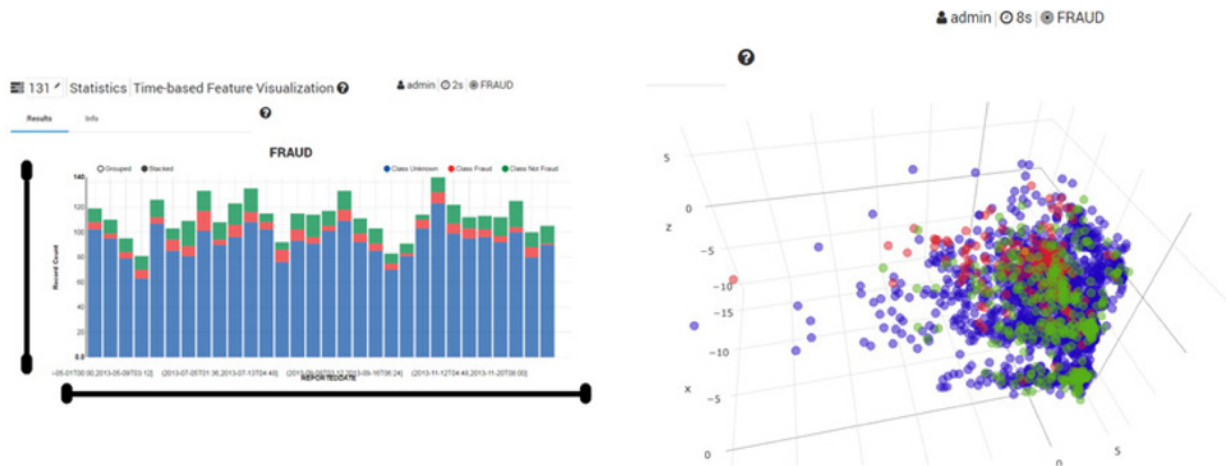
- AAP contains SymphonyAI NetReveal's big data engine for feature development, specifically to build enhanced datasets for modelling
- AAP exports its model outputs in the form of standard predictive modelling markup language (PMML) files which plug directly into the SymphonyAI NetReveal scoring manager

The full analytics ecosystem containing AAP and the SymphonyAI NetReveal production platform is shown below:



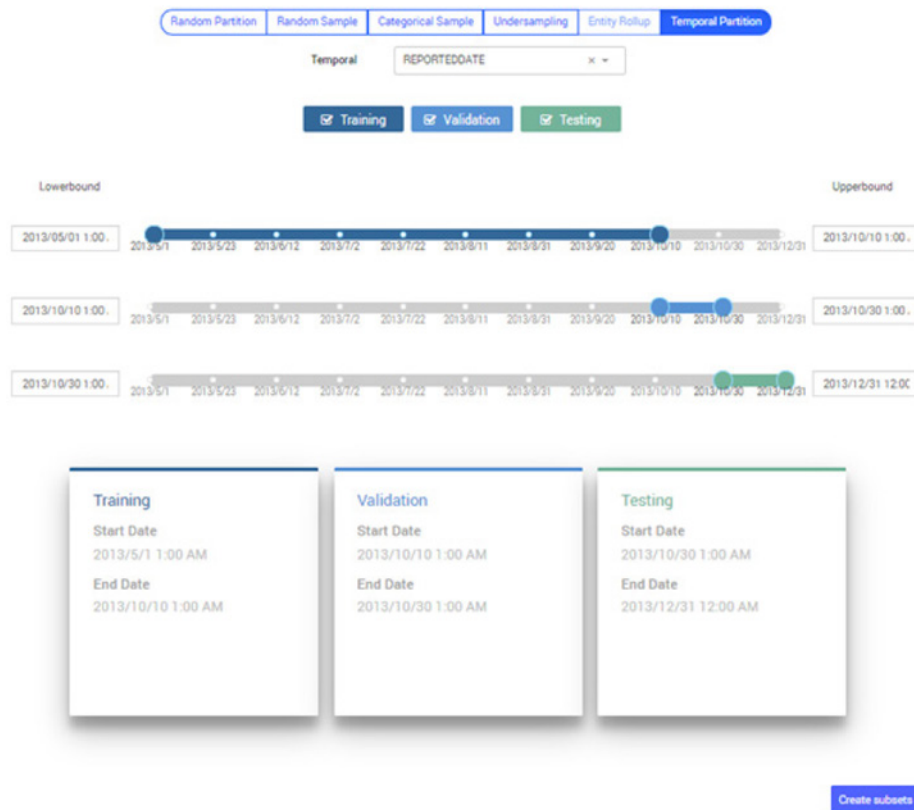
Data acquisition and feature development

AAP ingests flat files containing customer, account, and transaction data in addition to a processing model configuration file from SymphonyAI NetReveal that defines what fields are to be profiled and with respect to what time periods, etc. It is then straightforward within AAP to extract behavioural profile features from the raw data. AAP contains a range of statistical and visualisation-based functions that explore these features, in particular to remove correlated or low importance features. Illustrative screenshots from this part of the ecosystem are shown below:



Modelling and analytics

AAP contains a powerful set of machine learning analytics, both supervised and unsupervised models. In addition, AAP simplifies the crucial steps required to partition, sample, and normalise the data before these models are applied. For example, AAP employs a slider bar control for splitting data into training, validation, and test sets and a UI to filter the data based on a particular entity (e.g. account) or a particular customer attribute (e.g. segment). A screenshot is displayed below.



AAP's supervised machine learning analytics for classifying data into suspicious or not-suspicious classes include: rule induction, logistic regression, random forest, support vector machine, AdaBoost, gradient-boosted tree, and more. It is important to note that the non-suspicious class significantly overwhelms the suspicious class in typical datasets. This has to be carefully accounted for in any detection analytics. In AAP this is achieved by state-of-art under-sampling and oversampling techniques. Crucially, AAP relieves burden on the user associated with tuning the so-called hyper-parameters of its supervised machine learning models in order to achieve optimal detection performance. Specifically, AAP includes the option to auto-optimize these parameters.

AAP provides intuitive dashboard views that display the performance of multiple detection models and allow the user to rank them against a range of standard detection metrics, e.g. number of false positives, number of true positives, area under ROC curve, F1 score. This is shown in the screenshot below and the user may download this output as a png or csv file for inclusion in audit reports.

ID	Dataset	Featureset	Algorithm	Validation	ROC	Accuracy	TP%	FP%	FP:TP	F1	F2	F0.5	MCC	AUC	Parameters	Actions
18	May-Aug 2014-training-subset-8...	info50	Neural Networks	cross-5		92.4%	67.1%	4.2%	0.468	67.6%	67.3%	67.9%	0.633	0.938	Epochs: 200 Hidden Layer Size: 20 Learning Rate: -1 Momentum: 0.9 Scaling: Standardize Weight Decay: -6	
17	May-Aug 2014-training-subset-8...	info50	Neural Networks	cross-5		92.5%	68.2%	4.2%	0.461	68.3%	68.3%	68.4%	0.641	0.945	Epochs: 200 Hidden Layer Size: 10 Learning Rate: -1 Momentum: 0.9 Scaling: Standardize Weight Decay: -6	
16	May-Aug 2014-training-subset-8...	info50	Rule Induction	cross-5		92.7%	60.5%	3%	0.368	66.2%	62.6%	70.2%	0.625	0.784	Check Error Rate: true Grow/Prune Folds: 3 Optimization Runs: 2 Optimize Rules: true	
15	May-Aug 2014-training-subset-8...	info50	AdaBoost	cross-5		92.6%	61.4%	3.2%	0.385	66.4%	63.3%	69.7%	0.625	0.939	Leaf Nodes: 10 Num Trees: 50	
14	May-Aug 2014-training-subset-8...	Default	Logistic Regression	cross-5		91.3%	36.9%	1.5%	0.296	49.9%	41.2%	63.3%	0.495	0.822	Cost Minimizer: BFGS Iterations: 1000 Min Delta: 0.001 Regularization Coefficient: -7	

AAP also includes a drill-down view into any of the models to assess their individual performance in detail. An example is shown in the screenshot below. This shows a breakdown of the model's performance in terms of accuracy against detecting known suspicious instances (true positives) in a validation dataset and detecting known non-suspicious instances (true negatives) in the validation set. In any detection problem there is a trade-off between the number of true positives and the number of false positives. This trade-off is controlled by the decision boundary on the model's predictions and AAP exposes this to the user in the form a slider bar. AAP automatically optimises the decision boundary against a standard detection metric, but the user may override this value.

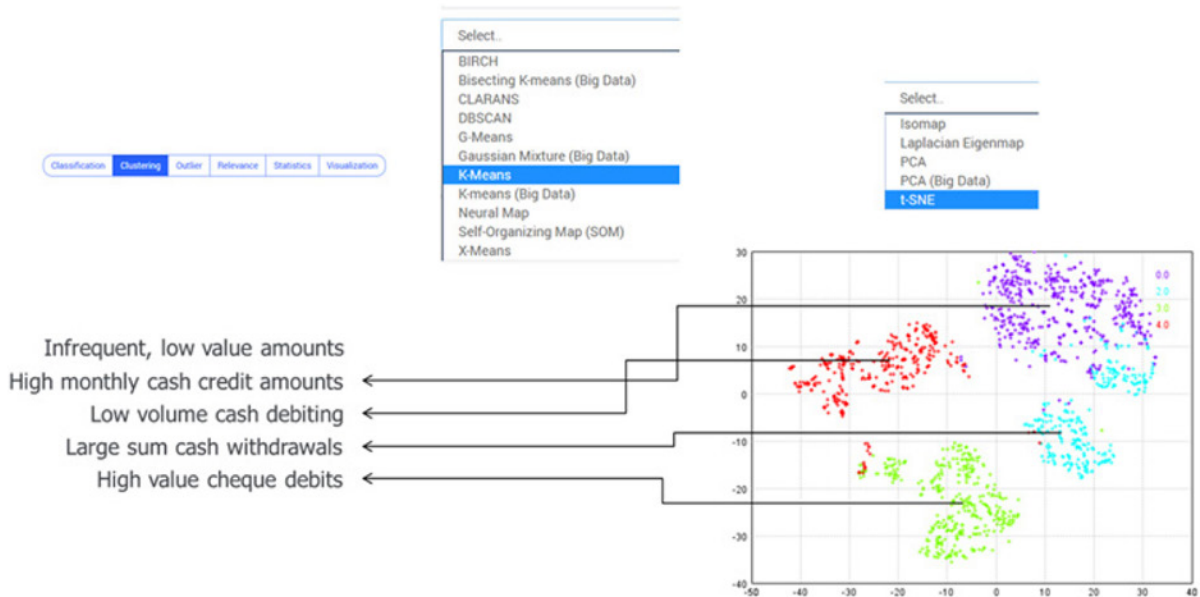


When AAP has built (or trained) a number of models they can be tested on fresh (or test) data. It is important to note that AAP generates white-box and auditable explanations for each suspicious event its machine learning models predict in the test dataset. This is highlighted in the screenshot below. Shown is a sample of predicted suspicious events and the features / feature values that the model has determined to be the most discriminating ones together with their importance weights.

Top Predictions					Reasons		
RowID	Key	Target Class	Probability	Reasons			
146	TXN-387-20090120	S_AML_AL_002	0.89	+	-		
177	TXN-237-20090121	S_AML_AL_002	0.87	+			
677	TXN151599	S_AML_AL_002	0.87	+			
254	TXN-208-20090123	S_AML_AL_002	0.87	+			
592	TXN151416	S_AML_AL_002	0.86	+			
					Feature	Value	Weight
					ExternalCreditMaximum max month present	3048041.26	17.49
					ExternalCreditAmount acc week present	3048041.26	17.43
					CashCreditAmount14 acc day total	3048041.26	16.27


AAP's unsupervised machine learning analytics for segmenting data into groups of clusters, e.g. customers who exhibit similar transactional patterns of behaviour, include: K-means, DBSCAN, Gaussian mixture model, self-organising Maps, and more. An example of a set of behavioural clusters that AAP discovered in a dataset is shown below. In AAP it is possible to not only determine the clusters (in this case using the K-means analytic) but also to visualise them via a 2-D embedding of the data (here using the t-SNE analytic).

AAP's unsupervised machine learning analytics for detecting outlier events in data, include: Isolation forest, K-nearest neighbours, local outlier factor, and local outlier probabilities. These analytics discover statistical anomalies in data and provide details with regard the underlying features / feature values that gave rise to each anomaly. The use case for these analytics is the so called "unknown unknowns" – patterns of financial crime that have not been seen before or even anticipated and therefore not represented by any of the existing detection scenarios.

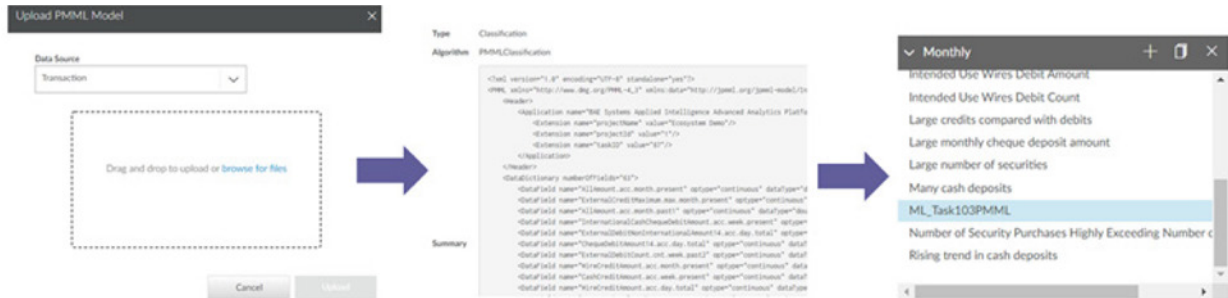


Testing, simulation, validation, and deployment

AAP enables one-click simulation-based testing of any of its trained predictive models alongside the detection scenarios contained within SymphonyAI NetReveal. This results in a further dashboard view comparing their respective performances with respect to true and false positives.



As indicated earlier, AAP's machine learning models are exported in the form of PMML files. These can be simply dragged-and-dropped into SymphonyAI NetReveal's user interface for deployment in the detection engine. This process is illustrated in the screenshot below:



An example of a user alert in SymphonyAI NetReveal that has been scored by an AAP machine learning model (Gradient Boosted Tree) is shown in the screenshot below:

Alert Identifier	Type/Sub-Type	Description	Age in days	Assigned To	Main Customer	Customer Segment	Customer Previously Reported	Case Name	Status	Main Employee Name	Organization Unit	Related Cases	Other Details
A2019112713201	AML - Suspicious Activity	Gradient Boosted Tree class = 1.0 Confidence = 0.8...	61	System User	Bob's Warehouse	51	No		Triage		North America		
A2019112713202	AML - Suspicious Activity	D503: Cash deposit = 21000.0 euro Threshold = \$1000...	61	Assign to me	Jack's Restaurant	51	No		Hibernated		North America		

Type a value	from	Type a value	Type a value	Type a value	Type a value	Type a value
to						
Time Period	Event Date	Transaction ID	Reason	Scenario Name	Account ID	Contributing Transactions
ForEachRecord	13/02/2016 00:00:00	UCD1_Txn_BW_1	Gradient Boosted Tree class = 1.0 Confidence = 0.8	AI - 1 : Gradient Boosted Tree	ACC_UCD1_BW_1	View

About SymphonyAI NetReveal

SymphonyAI NetReveal, part of the SymphonyAI finance vertical, is a global leader in regulatory compliance, fraud and financial crime detection, investigation, and reporting. The 750 SymphonyAI NetReveal digital, AI, and data experts work collaboratively with financial institutions across the globe to detect and prevent financial crime and optimize risk management operations.

Request a demo or contact us for more information:
netreveal.ai/request-demo