INTELLIGENCE · ARTIFICIAL AND COLLECTIVE INTELLIGENCE

Machine Learning for Market Intelligence

Machine learning lets market intelligence teams predict customer behaviour, group customers, and read large volumes of text. This article explains the main techniques, what they are good for, and the practical mistakes that make models less useful than they appear.

ModQu research article · 1,061 words · 11 references (APA 7th edition)

Market intelligence teams have always looked for patterns: which customers are likely to leave, which products sell together, what buyers complain about. **Machine learning** automates much of that search. Instead of an analyst writing rules, a model learns patterns from examples and applies them to new data. Used well, it lets a small team cover far more ground. Used badly, it produces confident numbers that mislead.

What machine learning does

Jordan and Mitchell (2015) describe machine learning as the study of systems that improve their performance on a task through experience, that is, through data. Two families matter most in market intelligence. **Supervised learning** learns to predict a known outcome, such as whether a customer cancelled or how much a store sold, from examples where the outcome is recorded. **Unsupervised learning** finds structure in data without a target, such as groups of similar customers or recurring topics in reviews.

Breiman (2001) drew an influential distinction between two cultures of statistical modelling. The traditional culture builds models to explain how data are generated. The algorithmic culture, which includes most machine learning, aims at accurate prediction and is less concerned with explanation. Market intelligence needs both: predictions to act on, and explanations to persuade managers and guide strategy.

Common tasks

Ma and Sun (2020) review the growing use of machine learning in marketing and show how it now supports most stages of market analysis. Figure 1 summarises the tasks most often used.

Machine learning tasks in market intelligence

Each task matches a type of question

Task	Question answered	Method type	Example data
Classification	Will this customer leave?	Supervised	Usage, billing history
Regression	How much will we sell?	Supervised	Sales, prices, season
Clustering	Which customers are alike?	Unsupervised	Purchases, survey answers
Topic modelling	What are people talking about?	Unsupervised	Reviews, forum posts
Anomaly detection	What has changed unusually?	Either	Prices, web traffic
Text classification	Is this comment a complaint?	Supervised	Calls, emails, reviews

Compiled from Ma and Sun (2020) and Jordan and Mitchell (2015); table by ModQu

ModQu.com

Figure 1. Machine learning tasks in market intelligence. Choosing the right task starts with being clear about the question. Compiled from Ma and Sun (2020) and Jordan and Mitchell (2015); table by ModQu.

Churn prediction is the classic example. Neslin et al. (2006) organised a tournament in which many teams built models to predict which mobile phone customers would leave. The best methods identified likely leavers far better than chance, and the differences between methods translated into meaningful sums of money. But Ascarza (2018) added an important caution: the customers most likely to leave are not always those most likely to respond to a retention offer. Targeting by risk alone can waste money. The useful question is whose behaviour the offer will change.

Text mining is the second major area. Netzer et al. (2012) showed that the way consumers mention brands together in online forums can reveal how a market is structured, recovering competitive maps similar to those from traditional surveys. Tirunillai and Tellis (2014) used topic modelling on product reviews to identify the dimensions of quality that consumers care about and how they change over time.

Practical mistakes

Domingos (2012) summarised lessons that practitioners often learn the hard way. Several apply directly to market intelligence:

1. **Overfitting.** A model that fits past data too closely captures noise and fails on new data. Always test on data the model has not seen.
2. **Leakage.** If information about the outcome slips into the inputs, for example a "cancellation date" field in a churn model, accuracy looks excellent and then collapses in practice.
3. **More data beats cleverer algorithms.** Better and more relevant data usually improve results more than a more sophisticated method.

4. **Correlation is not causation.** A model that predicts sales from advertising does not show that more advertising will raise sales.

Hastie, Tibshirani and Friedman (2009) provide the standard technical treatment of these issues, including methods for testing models fairly and choosing between them.

Two further problems arise after a model is in use. Markets change, and the patterns a model learned may stop holding. Gama et al. (2014) call this **concept drift** and describe methods for detecting it; the simplest is to monitor accuracy continuously and retrain when it falls. Sculley et al. (2015) warn that machine learning systems accumulate hidden **technical debt**: dependencies on data sources, undocumented assumptions and fragile pipelines that make them costly to maintain. A model is not a one-off project but an ongoing commitment.

An illustration

A subscription software company wants to reduce cancellations. Its intelligence team builds a classification model from two years of usage data. The model flags customers whose logins have fallen and who have opened support tickets. Rather than sending discounts to everyone flagged, the team runs a small trial: half the flagged customers receive a call from an account manager and half do not. The calls reduce cancellations among mid-sized accounts but make no difference among the smallest. The company focuses its effort accordingly. The model found the risk; the experiment showed where action paid off.

Explaining results to decision-makers

A model's output is only useful if managers trust it enough to act and understand it well enough to act sensibly. Three practices help. First, show which inputs drive the predictions, in plain language: "customers whose logins fell by half in a month are three times as likely to cancel". Second, compare the model with the current way of deciding, so that the gain is concrete. Third, be candid about where the model is weak, for example with new customers who have little history. Breiman (2001) noted that the most accurate models are often the hardest to interpret; where a decision affects many people or large sums, a slightly less accurate model that managers understand may be the better choice.

Combining models with other intelligence

Machine learning works best alongside other sources. Survey data can explain why customers behave as the model predicts. Interviews with sales staff can reveal changes the data have not yet captured. Competitor monitoring can warn that a pattern is about to break. The model provides scale and consistency; the other sources provide context.

Getting started

For teams new to machine learning, a few habits help. Start with a clear business question and a measurable outcome. Use simple, interpretable models first and move to complex ones only if they clearly perform better. Keep a hold-out sample for honest testing. Document the data, assumptions and known weaknesses of each model. And always present results with the uncertainty attached: a prediction that a customer has a 70% chance of leaving is more useful, and more honest, than a label saying "will churn".

Conclusion

Machine learning extends what market intelligence teams can see, from individual customer risk to the themes in millions of comments. Its value depends on asking the right question, testing honestly, watching for change and combining prediction with judgement about what action will actually work.

References

Ascarza, E. (2018). Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing Research*, *55*(1), 80–98. https://doi.org/10.1509/jmr.16.0163

Breiman, L. (2001). Statistical modeling: The two cultures. *Statistical Science*, *16*(3), 199–231. https://doi.org/10.1214/ss/1009213726

Domingos, P. (2012). A few useful things to know about machine learning. *Communications of the ACM*, *55*(10), 78–87. https://doi.org/10.1145/2347736.2347755

Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, *46*(4), Article 44. https://doi.org/10.1145/2523813

Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Springer. https://doi.org/10.1007/978-0-387-84858-7

Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, *349*(6245), 255–260. https://doi.org/10.1126/science.aaa8415

Ma, L., & Sun, B. (2020). Machine learning and AI in marketing – Connecting computing power to human insights. *International Journal of Research in Marketing*, *37*(3), 481–504. https://doi.org/10.1016/j.ijresmar.2020.04.005

Neslin, S. A., Gupta, S., Kamakura, W., Lu, J., & Mason, C. H. (2006). Defection detection: Measuring and understanding the predictive accuracy of customer churn models. *Journal of Marketing Research*, *43*(2), 204–211. https://doi.org/10.1509/jmkr.43.2.204

Netzer, O., Feldman, R., Goldenberg, J., & Fresko, M. (2012). Mine your own business: Market-structure surveillance through text mining. *Marketing Science, 31*(3), 521–543. https://doi.org/10.1287/mksc.1120.0713

Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. In *Advances in Neural Information Processing Systems 28* (pp. 2503–2511).

Tirunillai, S., & Tellis, G. J. (2014). Mining marketing meaning from online chatter: Strategic brand analysis of big data using latent Dirichlet allocation. *Journal of Marketing Research, 51*(4), 463–479. https://doi.org/10.1509/jmr.12.0106