INTELLIGENCE · ANALYSIS AND TRADECRAFT

Activity-Based Intelligence and Pattern-of-Life Knowledge Generation

Activity-based intelligence studies the activities and transactions of people and organisations over time and place, rather than starting from a known target. This article explains its principles, the science of human routine on which it rests, its uses in security, policing and business, and the serious privacy and fairness questions it raises.

ModQu research article · 2,430 words · 23 references (APA 7th edition)

Traditional intelligence starts with a target. An analyst knows which organisation, facility or person matters and directs collection towards it. But some of the most important intelligence problems begin without a known target. Who is building the network of safe houses? Which suppliers are quietly serving a sanctioned company? Where will the next surge in demand appear? **Activity-based intelligence** (ABI) addresses such questions by studying activities and transactions across time and space, looking for patterns and for breaks in patterns, and discovering targets from the data rather than starting with them. Closely related is **pattern-of-life analysis**: building a picture of the normal routines of a place, group or person, so that unusual activity stands out.

This article explains the principles of ABI, the research on human routine that makes it possible, its applications in security, policing and business, and the privacy, fairness and accuracy concerns it raises.

Origins and principles

ABI developed in American military and intelligence organisations during the conflicts in Iraq and Afghanistan in the 2000s. Adversaries there did not operate from fixed military installations; they hid among civilian populations and were identified mainly through their activities: meetings, movements, communications and transactions. At the same time, the volume of data from aerial surveillance, signals collection and other sensors grew far beyond what analysts could examine one report at a time.

Atwood (2015) and Biltgen and Ryan (2016) describe four principles that distinguish ABI from traditional, target-first intelligence.

Four principles of activity-based intelligence

How ABI differs from target-first collection

Principle	What it means	Why it matters
Georeference to discover	Tag every item with place and time	Links appear between data that share a location
Data neutrality	Treat all sources as potentially useful	The key clue may come from an unexpected source
Sequence neutrality	Data collected before a question can answer it	Archives allow the past to be re-examined
Integrate before exploiting	Combine data before deep analysis	Patterns show only when sources are joined

Based on Atwood (2015) and Biltgen and Ryan (2016); table by ModQu

ModQu.com

Figure 1. Four principles of activity-based intelligence. Together they shift the emphasis from collecting on known targets to discovering unknown ones in combined data. Based on Atwood (2015) and Biltgen and Ryan (2016); table by ModQu.

Georeference to discover means tagging every piece of data with where and when it was produced, because things that happen at the same place and time are often connected. **Data neutrality** means treating all sources as potentially valuable, since the critical clue may come from an unexpected one. **Sequence neutrality** recognises that data collected before a question was asked may later answer it, which makes archives valuable. **Integration before exploitation** means combining data from different sources before detailed analysis, because many patterns become visible only when sources are joined.

ABI fits within the broader discipline of geospatial intelligence described by Lowenthal and Clark (2016), but it is less a collection discipline than an analytic method. It complements the target-centric approach advocated by Clark (2020): ABI discovers targets, which can then be modelled and studied in depth.

The science of human routine

ABI depends on a basic fact about human behaviour: it is highly regular. Geographers recognised this long before modern data existed. Hägerstrand (1970) developed **time geography**, which traces each person's path through space and time and shows how their activities are constrained by the need to be in certain places at certain times, by travel time and by the opening hours of the places they visit.

Large data sets have since confirmed and quantified this regularity. González et al. (2008) analysed the movements of 100,000 mobile phone users over six months and found that, despite wide differences in how far people travelled, individual movement followed simple, reproducible patterns, with most people returning repeatedly to a few locations. Song et al. (2010) estimated that, in principle, a person's location could be predicted correctly most of the time from their past movements, because human mobility is so regular. Eagle and Pentland (2006), in their "reality mining" study, used data from the phones of about a hundred people to infer their daily routines and social relationships, and showed that deviations from routine could be detected automatically.

Such data can reveal more than movement. Blumenstock et al. (2015) showed that patterns of mobile phone use in Rwanda could be used to estimate the wealth of individuals and the distribution of poverty across the country. Kitchin (2014) describes how cities are increasingly instrumented with sensors that record traffic, energy use and movement in real time, creating a continuous record of urban activity.

Finding anomalies

If normal activity is regular, unusual activity can stand out. The core analytic task of ABI is to establish what is normal and then find departures from it.

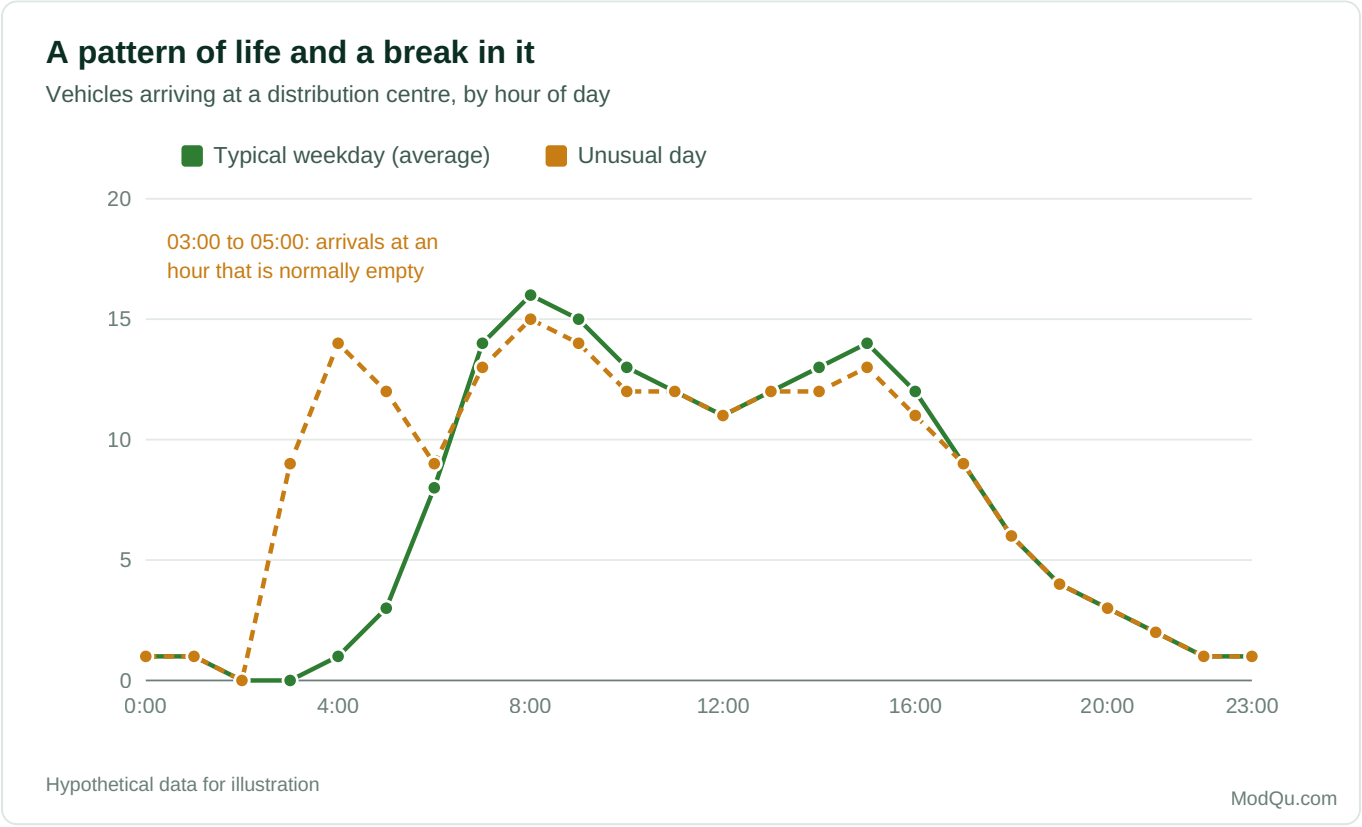


Figure 2. A pattern of life and a break in it. The unusual day differs from the norm only in the early hours, which is exactly the kind of change pattern-of-life analysis is designed to catch. Hypothetical data for illustration.

Chandola et al. (2009), in a widely cited survey of anomaly detection methods, distinguish three kinds of anomaly. A **point anomaly** is a single observation that is unusual in itself, such as a vehicle at a site that is normally empty. A **contextual anomaly** is normal in one context but not another: heavy traffic at a warehouse is normal at midday but unusual at three in the morning. A **collective anomaly** is a group of observations that is unusual together, though each might be normal alone, such as several unrelated companies registering at the same address within a week.

Not every anomaly is significant. Most have innocent explanations: a public holiday, a new delivery schedule, a data error. Heuer (1999) warned that analysts tend to see patterns and causes even in random data. Anomalies should therefore be treated as prompts for investigation, not as conclusions.

Building a baseline

Everything in pattern-of-life analysis depends on a good description of what is normal. Building that baseline takes more care than it might appear. Activity varies by hour, by day of the week, by season and around holidays, so a single average can hide the very variation that matters. A warehouse that is quiet on Sunday mornings is not anomalous on a Sunday; the same quiet on a Tuesday may be. Analysts therefore usually build baselines for each relevant context, such as each hour of each weekday, and describe not only the average level of activity but also its usual range.

Baselines also change. A new shift pattern, a change of ownership or a new road can alter the normal pattern permanently. If the baseline is not updated, every subsequent day will look anomalous and the analysis will drown in false alarms. If it is updated too quickly, a genuine change may be absorbed into the new normal before anyone notices it. Biltgen and Ryan (2016) emphasise that maintaining the baseline is a continuing analytic task, not a one-off calculation.

Finally, the baseline is only as good as the data behind it. Gaps in coverage, changes in sensors and errors in timestamps can all create apparent patterns. Before interpreting a change in activity, analysts should check whether the way the activity was recorded has changed.

The problem of false alarms

Anomaly detection faces a basic statistical difficulty. When genuinely significant events are rare, even an accurate detection method will produce many more false alarms than true ones. Suppose a method flags 1 per cent of ordinary days as unusual, and a significant event occurs on one day in a thousand. Over a thousand days, the method will flag about ten ordinary days and, at best, one significant one. Most alerts will therefore be false, however good the method.

This has practical consequences. Analysts need a way to prioritise alerts, for example by combining several weak signals before acting, or by focusing on anomalies that appear across more than one data source. Organisations need to accept that investigating false alarms is part of the cost of the method. And decision-makers should be told how often past alerts proved significant, so that they can judge how much weight to give a new one.

An ABI workflow

In practice, activity-based analysis tends to follow a recurring sequence. Analysts first assemble and georeference the available data, joining sources by place and time. They then describe the baseline pattern for the places, networks or entities of interest. Next, they search for anomalies and for unexpected connections, such as the same vehicle, account or address appearing in several otherwise separate activities. Promising leads are investigated with additional data and, where appropriate, traditional collection. Finally, what has been learned is fed back: newly identified entities become targets for closer study, and the baseline and detection rules are refined. Clark (2020) describes a similar cycle of discovery and modelling, in which each round of analysis improves the model of the target and suggests what to look for next.

Lessons from policing

Policing has used pattern analysis of place and time for decades and offers useful lessons for ABI more broadly. Cohen and Felson (1979) proposed **routine activity theory**: crime occurs when a motivated offender and a suitable target come together in the absence of a capable guardian, and this convergence depends on people's everyday routines. Sherman et al. (1989) found that in Minneapolis a very small proportion of addresses produced about half of all calls to the police, and called these places **hot spots**. Spatial statistics, such as the local indicators of spatial association developed by Anselin (1995), allow analysts to test whether such clusters are real or could have arisen by chance.

Braga et al. (2014), in a systematic review and meta-analysis, found that focusing police attention on hot spots reduces crime, generally without simply displacing it to nearby areas. Ratcliffe (2016) describes how such analysis became central to **intelligence-led policing**, in which data analysis guides where and how police resources are used.

Predictive policing extended these ideas with statistical models that forecast where crimes are likely to occur. Mohler et al. (2015) conducted field trials in which a model's forecasts outperformed those of experienced crime analysts in identifying locations for patrols. But Lum and Isaac (2016) demonstrated a serious risk. Using data on drug arrests in Oakland, California, they showed that a predictive model would direct police disproportionately to minority neighbourhoods, because the arrest data reflected past policing patterns rather than the true distribution of drug use. More patrols would produce more arrests in those areas, which would feed back into the model. The lesson applies to all ABI: models learn from the data they are given, including its biases.

Applications in business

The same methods are widely used in business, usually under other names.

- **Fraud detection.** Banks and payment companies model each customer's normal pattern of transactions and flag departures, such as a purchase in an unusual place or at an unusual time.
- **Supply chain monitoring.** Shipping, port and logistics data reveal patterns of trade, delays and changes in suppliers.
- **Retail and location analysis.** Data on footfall and movement show how customers use stores and neighbourhoods, and how patterns change after a competitor opens nearby.

- **Competitive intelligence.** Activity at a competitor's facilities, such as traffic, hiring and shipping, can signal changes in production before they are announced.
- **Operational monitoring.** Patterns in equipment sensor data reveal developing faults before they cause breakdowns.

In each case, the method is the same: establish the normal pattern, detect departures, and investigate whether they matter.

Privacy and ethics

ABI raises serious privacy concerns, because the data that reveal patterns of life are, by definition, data about people's lives. Mayer et al. (2016) studied telephone metadata, the records of who called whom and when, without the content of calls. They found that such metadata could reveal sensitive information, including medical conditions, religious practice and relationships, and could often be linked to identifiable individuals.

Solove (2006) developed a taxonomy of privacy harms that includes **aggregation**: combining pieces of information that are harmless individually into a profile that reveals far more than any one piece. ABI is, in essence, an aggregation technique, which means that its privacy implications can be serious even when each data source seems innocuous. Nissenbaum (2004) proposed the idea of **contextual integrity**: information flows are appropriate when they respect the norms of the context in which the information was shared. Location data shared with a navigation app to find a route carry an expectation that they will be used for that purpose, not to build a profile of the user's daily life.

Zuboff (2019) argues that the commercial collection of behavioural data has created a new economic order, in which companies profit from predicting and shaping behaviour, often without meaningful consent. Whatever one thinks of that argument, it highlights the importance of legitimacy. Organisations using pattern-of-life methods should ask not only whether the analysis is legal but whether the people whose activities are studied would consider it fair.

Practices for responsible ABI

1. **Be clear about purpose.** Define what question the analysis serves and limit data collection and retention to what that purpose requires.
2. **Prefer aggregate patterns.** Many business questions can be answered with patterns across groups, without identifying individuals.
3. **Respect context and consent.** Use personal data in ways consistent with how and why it was collected, and comply with privacy law.
4. **Check the data for bias.** Ask how the data were generated and whether they reflect the phenomenon of interest or the way it has been observed (Lum & Isaac, 2016).
5. **Treat anomalies as questions.** Investigate before acting, and consider innocent explanations.
6. **Keep human judgement.** Automated detection should support analysts, not replace their assessment of what a pattern means.

An illustration

A Canadian logistics company notices that delivery times in one region have become erratic. Rather than investigating each delay separately, its analysts combine data from vehicle tracking, warehouse scanners, weather records and road works notices, all referenced to place and time. A pattern emerges: delays cluster on weekday mornings along one corridor, beginning on the date a bridge was partly closed for repairs. Further analysis shows that drivers have shifted to a route through a residential area where schools open at the same time as the morning dispatch. The company changes its dispatch schedule in that region, and delivery times return to normal within two weeks. No individual driver's data were needed beyond the vehicle level, and the analysis answered a clear operational question.

Conclusion

Activity-based intelligence turns the traditional approach around: instead of collecting on known targets, it studies activity across place and time to discover what matters. It rests on the strong regularity of human behaviour, which large data sets have repeatedly confirmed, and on methods for detecting departures from normal patterns. Policing research shows both its power, in focusing attention where it is most needed, and its risks, when biased data produce biased predictions. Because patterns of life are patterns of people's lives, ABI also raises serious privacy questions. Used with a clear purpose, attention to data quality and respect for privacy, it is a powerful method for generating knowledge from the record of what people and organisations actually do.

References

Anselin, L. (1995). Local indicators of spatial association—LISA. *Geographical Analysis, 27*(2), 93–115. https://doi.org/10.1111/j.1538-4632.1995.tb00338.x

Atwood, C. P. (2015). Activity-based intelligence: Revolutionizing military intelligence analysis. *Joint Force Quarterly, 77*, 24–33.

Biltgen, P., & Ryan, S. (2016). *Activity-based intelligence: Principles and applications*. Artech House.

Blumenstock, J., Cadamuro, G., & On, R. (2015). Predicting poverty and wealth from mobile phone metadata. *Science, 350*(6264), 1073–1076. https://doi.org/10.1126/science.aac4420

Braga, A. A., Papachristos, A. V., & Hureau, D. M. (2014). The effects of hot spots policing on crime: An updated systematic review and meta-analysis. *Justice Quarterly, 31*(4), 633–663. https://doi.org/10.1080/07418825.2012.673632

Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys, 41*(3), Article 15. https://doi.org/10.1145/1541880.1541882

Clark, R. M. (2020). *Intelligence analysis: A target-centric approach* (6th ed.). CQ Press.

Cohen, L. E., & Felson, M. (1979). Social change and crime rate trends: A routine activity approach. *American Sociological Review, 44*(4), 588–608. https://doi.org/10.2307/2094589

Eagle, N., & Pentland, A. (2006). Reality mining: Sensing complex social systems. *Personal and Ubiquitous Computing*, *10*(4), 255–268. https://doi.org/10.1007/s00779-005-0046-3

González, M. C., Hidalgo, C. A., & Barabási, A.-L. (2008). Understanding individual human mobility patterns. *Nature*, *453*(7196), 779–782. https://doi.org/10.1038/nature06958

Heuer, R. J., Jr. (1999). *Psychology of intelligence analysis*. Center for the Study of Intelligence, Central Intelligence Agency.

Hägerstrand, T. (1970). What about people in regional science? *Papers of the Regional Science Association*, *24*(1), 7–24. https://doi.org/10.1007/BF01936872

Kitchin, R. (2014). The real-time city? Big data and smart urbanism. *GeoJournal*, *79*(1), 1–14. https://doi.org/10.1007/s10708-013-9516-8

Lowenthal, M. M., & Clark, R. M. (Eds.). (2016). *The five disciplines of intelligence collection*. CQ Press.

Lum, K., & Isaac, W. (2016). To predict and serve? *Significance*, *13*(5), 14–19. https://doi.org/10.1111/j.1740-9713.2016.00960.x

Mayer, J., Mutchler, P., & Mitchell, J. C. (2016). Evaluating the privacy properties of telephone metadata. *Proceedings of the National Academy of Sciences*, *113*(20), 5536–5541. https://doi.org/10.1073/pnas.1508081113

Mohler, G. O., Short, M. B., Malinowski, S., Johnson, M., Tita, G. E., Bertozzi, A. L., & Brantingham, P. J. (2015). Randomized controlled field trials of predictive policing. *Journal of the American Statistical Association*, *110*(512), 1399–1411. https://doi.org/10.1080/01621459.2015.1077710

Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, *79*(1), 119–157.

Ratcliffe, J. H. (2016). *Intelligence-led policing* (2nd ed.). Routledge.

Sherman, L. W., Gartin, P. R., & Buerger, M. E. (1989). Hot spots of predatory crime: Routine activities and the criminology of place. *Criminology*, *27*(1), 27–56. https://doi.org/10.1111/j.1745-9125.1989.tb00862.x

Solove, D. J. (2006). A taxonomy of privacy. *University of Pennsylvania Law Review*, *154*(3), 477–564. https://doi.org/10.2307/40041279

Song, C., Qu, Z., Blumm, N., & Barabási, A.-L. (2010). Limits of predictability in human mobility. *Science*, *327*(5968), 1018–1021. https://doi.org/10.1126/science.1177170

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.