



High Income,
Lives in Small Town

Large Family

Rare Disease

CONFIDENTIAL OPEN DATA

Safe Data Through Disclosure Control



Government departments and agencies around the world are facing increasing pressure to provide more access to data they have collected. This pressure is accompanied by demands that data access be interactive and meet the needs of end users, rather than the reporting requirements of the organisation releasing the data.

This poses a huge dilemma for modern organisations: how can they balance the increasing call for data with their legal and duty-of-care requirements to protect the privacy of individuals that have contributed to that data.

This white paper looks at the challenges and solutions currently available for confidentiality protection and provides a process for you to discover the right level of protection for your needs.

The Case for Disclosure Control

A breach of confidentiality occurs when a person or entity is recognised in a dataset, allowing an attacker to find out new information about that person.

Yet confidentiality is about more than just removing personally identifiable information (PII) from your data. While removing independently sensitive fields such as names, credit card numbers or IP addresses is a minimum level of protection, this alone will not always hide the identity of the individuals in the data. This is especially true with web access interfaces that allow multiple queries to be submitted and easy ways to digitally combine datasets.

True privacy protection requires an integrated and systematic approach. Disclosure control is the name given to this system of hardware, software and statistical solutions that combine to create a safe environment for data dissemination.

Rights and Obligations

Data providers must maintain the privacy of the individuals and organisations that contribute to their data, and many countries have laws and regulations that require and reinforce this.

It is also important to note that malicious identity theft is not the only consequence of confidentiality breaches. If respondents do not believe they are adequately protected from a possible information threat, they are less likely to comply with requests for information, or submit the correct details. For organisations involved in planning, this can have far reaching consequences.

On the other hand, if an organisation provides data that has been confidentialised to such an extent that it is less

“

The government's open data agenda allows us to find out more than ever about the performance of public bodies.

However, there is also a risk that we will be able to piece together a picture of individuals' private lives.

With ever increasing amounts of personal information in the public domain, it is important that organisations have a structured and methodical approach to assessing the risks.¹

**Christopher Graham
UK Information
Commissioner**

useful and accessible, then it faces negative reactions from end users. This can also increase staff workload, as they struggle to serve a growing number of more complex data queries from consumers and stakeholders. For the modern data collection agency, balancing this risk–utility equation is one of the toughest tasks in designing disclosure control.

How is Confidentiality Breached?

A breach of confidentiality can happen when multiple fields or variables are combined to uniquely identify an individual or enterprise. While “occupation”, “postcode” and “number of dependents” might not on their own be considered identifying information, what if there is only one person in the dataset that matches a specific combination of all three?

Typical confidentiality breaches involve combining information in the released data with some known or easily discoverable information. For example:

- An attacker searches for a specific individual in the dataset, based on information the attacker already knows about that individual. Only one record matches all known criteria.
- An attacker starts with a record in the anonymised dataset and then tries to identify that individual by matching them with publicly available information.

Some typical types of attacks are described below.

Differencing Attacks

A differencing attack involves generating two related tables and comparing the results. For example, one for all employees and one for employees earning less than \$150,000. By subtracting the results of the two tables an

A 2014 study of anonymised credit card data belonging to 1.1 million people found that just four pieces of external information were enough to match a person with their anonymised credit card record, 90% of the time.²

attacker can produce a third, “differenced” table containing information about a subset of interest (in this case, employees earning *over* \$150,000).

Differencing attacks are commonly used to breach suppression algorithms, which are rules that hide cells with a count below a certain value or with fewer than a certain number of contributors. In a differencing attack, the tables created have results that are large enough not to trigger the suppression rules, allowing the attacker to infer results that would normally be suppressed.

Homogeneity Breaches

Sometimes confidentiality breaches occur without even needing to isolate a single record in the data. Consider a simple example table showing medical conditions for males employed by a particular organisation.

In this case, even though there is a count value of 5 for the embarrassing condition, all males in the selected subgroup have that condition. This group’s confidentiality has been breached through homogeneity of their attributes.

Department	Sales
Gender	Male
Age Band	50-54
Diabetes	0
Something Embarrassing	5
High Blood Pressure	0

Inferred Bounding

Inferred bounding is the name given to a targeted attack that attempts to calculate the value of a suppressed cell to within set “bounds”. This typically involves sophisticated linear programming techniques to deduce the maximum and minimum possible values of a protected cell. If the difference between the upper and lower bound is less than one count, then the suppressed value is discovered.

Although bounding can commonly be used on count datasets to deduce the count with a particular combination of attributes, inference can also successfully be used to deduce magnitude data.

Consider two competing firms that dominate industry in particular region. A simple query regarding the monetary value of government grants to each region immediately tells each firm roughly what its competitor is receiving.

Protecting against this risk is not as simple as enforcing a minimum number of contributors (say, three) to each cell. In the above scenario, there could easily be three firms, but one is much smaller than the others, in which case a good bounding estimate can still be generated from the report. This type of attack is prevalent in commercially sensitive information, particularly financial or otherwise tactically sensitive magnitude data.

Disclosure Control Solutions

When considering an effective disclosure control solution, many factors need to be taken into account, including the type of data, how the data is reported, the end users and the overall sensitivity of the dataset. Another consideration is where to apply disclosure control: at the microdata level or post query?

Microdata Confidentialisation

Microdata confidentiality involves pre-processing individual unit records, and is most suited to organisations that need to release microdata cubes (such as those with a high proportion of research users).




Typical methods include attribute swapping, generalisation techniques or tuple suppression, and involve finding suitable matches within the record set that need to be

adjusted. However, the number of records and classifiers in modern sets makes this a challenging task, one that can rapidly outclass realistic computational abilities.

Great care also needs to be taken to ensure that the overall statistical properties of the data are not compromised. Any bias or modified variance in the data must be kept to a minimum.

Pre-Aggregation


This technique involves pre-aggregating the unit records in some way, to protect access to the microdata. Business rules are applied to generate a “safe” level of aggregation below which the query engines may not drill.



Further disclosure control is applied to predefined output datasets, which are then released.

Query-Based Tabular Level Confidentialisation

In the examples above, end users can submit queries against predefined, approved output. A variant on this approach is query-based access, which relies upon the routines being sufficiently robust that any tabulation query can be submitted by an end user. This approach reduces and sometimes replaces the need to predefine output, delivering greater flexibility to the user.



Query-based access requires a tabular confidentiality routine that can work for a limitless number of *ad hoc* queries. This scenario removes the need for the data provider to predict what combinations of variables need to be combined in summary outputs, delivering much greater flexibility to the end user.

Solutions for Data Dissemination

Current disclosure control techniques fall into three categories:

Suppression

The most visible method is simply to hide sensitive values, replacing them with a symbol or a 0 value. However, a high level of protection can only be achieved when both sensitive cells and related cells are suppressed.

For example, if one value in a row is sensitive and is suppressed, then other values in the row, or the total for that row, may also need to be suppressed otherwise attackers may be able to calculate the suppressed value.

This can lead to very low utility, as many values in a table may be rendered unavailable. Suppression methods are also subject to breach through bounding or differencing attacks.

Generalisation

This method suppresses some values to a degree, by only reporting more general values (for instance, reporting a more general 4 digit region code, rather than a 5 digit postal code).

It can be applied at both the microdata or at the tabular

level, depending on the individual requirements, and is typically used in conjunction with other confidentiality mechanisms.

Obfuscation

Obfuscation techniques hide information by adjusting the true value of any given cell in the table and reporting a slightly different value instead. With these techniques it is important to ensure that the resulting table preserves the same statistical characteristics as the original table, and that no bias is introduced by changing values.

Some obfuscation techniques, such as randomly rounding cell values, offer high utility and are fast and easy to implement, but may reduce the usefulness of the data, if the rounding is too aggressive. This approach also may introduce a high level of bias, as the rounding is not controlled.

It is also possible that different users may see different results for the same query.

Perturbation

WingArc Australia's perturbation algorithm is a form of obfuscation. It makes adjustments to cell values to ensure that individuals cannot be identified. However, these adjustments are both controlled and repeatable.

This offers a good balance between utility and protection, and ensures that no bias is introduced. It also ensures that the same cell is always adjusted in exactly the same way, no matter how the table query is constructed.

It supports both count and magnitude data and has become a natural choice for many high risk datasets, including population census data.

“

WingArc Australia's perturbation algorithm enables the Australian government to allow safe online access to census microdata.³

Choosing the Right Solution for Your Data

Disclosure control methods are developing rapidly to accommodate the increasing demand for online data. The benefit of each approach must be weighed against the potential cost to data accessibility and utility.

To do this, we recommend that a data confidentiality assessment be carried out. This process involves developing a risk profile before data is released that takes into account factors such as:

- Data collection factors, such as whether the data comes from a census or survey, and any risks from survey frame or post-collection weighting.
- The type of data (count or magnitude).
- The size of the dataset.
- The range of end users and the type of access required.
- The likelihood and consequences of breaches, considering the end users, the sensitivity of the data, the existence of previous or similar datasets and the required level of detail in the data.
- Implementation and usability requirements.

There are a range of tabular and microdata confidentiality methods available for implementation. Some are already in regular use internationally, whilst new methods are being proposed and tested constantly.

A confidentiality assessment process can help organisations to:

- Build a customised risk profile for your data.
- Suggest the most appropriate disclosure control solution.
- Understand and communicate the benefits and risks of the preferred solution.

More Information

WingArc Australia works with government agencies and statistical offices around the world to design, build and test robust, fit-for-purpose confidentiality solutions. Our technology helps those agencies publish data safely and securely.

Our powerful perturbation technology provides an off-the-shelf solution for confidentiality, and our Disclosure Control API provides a very flexible platform for implementing a fully customised confidentiality solution for each customer.

Get in touch today to learn more about what we can do for you: <https://wingarc.com.au/contact/>

References

1. <https://ico.org.uk/media/1061/anonymisation-code.pdf>
2. "Unique in the shopping mall: On the reidentifiability of credit card metadata", *Science Magazine*, 30 Jan 2015 (Vol. 347, Issue 6221, pp. 536-539).
<http://science.sciencemag.org/content/347/6221/536>
3. <http://www.abs.gov.au/websitedbs/censushome.nsf/home/tablebuilder>

Company Overview

- > Incorporated as Space-Time Research in 1986 at Melbourne University to commercialise research into large database manipulation and interrogation
- > Acquired by WingArc1st, Japan's leader in Business Intelligence and Data Analytics, in 2017
- > WingArc Australia Headquarters in Melbourne
- > Global customer base

Solution Overview

- > Self-service data analytics
- > Data dissemination
- > Easy to use: anyone can analyse big data sets
- > Advanced confidentiality and privacy protection
- > Advanced statistical capabilities
- > High speed, high performance tabulation engine



WingArc Australia

360 Collins Street
Melbourne
Victoria, 3000
Australia

+61 3 9615 5200

wingarc.com.au