

Temporal Rules and Queries In the Antimicrobial-Resistance Domain: Collection of Examples

Jing Ao¹ and Rada Chirkova²

^{1,2}Department of Computer Science, North Carolina State
University

1 Introduction

In this report we provide several examples that focus on answering temporal questions on temporal data, in contexts relevant to the antimicrobial-resistance (AMR) domain, see [7] for a discussion of some AMR scenarios of interest. We consider the *data-exchange setting*, in which the data of interest are stored in a *source* database, and the queries expressing the relevant questions are posed on the schema of the *target* database assumed to not contain any data. The *data-exchange* process enables evaluation of the target queries, via populating the target database by importing the source data into it, using *source-to-target dependencies* (*s-t tgds*). (For the details, please see, e.g., [4].) This report uses the data-exchange definitions and notation as provided in [4].

The six examples provided in this report, see Section 2, focus on *temporal relational-to-relational data exchange and query answering*. In this setting, the target temporal data are stored as relational tables that use additional temporal attributes to capture the temporal information from the source relations [3]. Further, the temporal questions of interest are specified as SQL queries. Examples 1–4 use GAV mappings [5] as the data-exchange rules. Among them, Examples 1 and 2 deal with real-world AMR data collected at laboratories in the spreadsheet format, and Example 3 leverages real-world aggregated AMR data. In the report, we use real-world schemata for these data, while making up some of the data values for the purpose of making the exhibited s-t tgds “go through.” Although Example 4 uses made-up schemata and partially made-up data, the meaning of each attribute and the relationships between the attributes are reasonable to real-world AMR domain experts, as confirmed by our collaborators on the AMR team led by Dr. Paula Cray at NCSU. Examples 5 and 6 leverage GLAV source-to-target mappings.

2 Temporal Relational-to-Relational Data Exchange and Query Answering

2.1 Example 1

2.1.1 Source Schema and Instance

This example uses real AMR data shared by the research team led at NCSU by Dr. Paula Cray, see Table 1. (The attribute *Sal ID* refers to the isolate of a kind of bacteria specified by *Serotype*, and *TET MIC* represents the resistance-testing reading of this bacteria with respect to the AMR drug *TET*, collected on animals.)

AnimalReadings				
Date	Farm Name	Sal ID	Serotype	TET MIC
[5/19/09, 5/19/09]	Louisburg Farm	S277	S.Anatum	>32
[8/27/17, 8/27/17]	Houston Farm	S2020	S.Anatum	>32

Table 1: Resistance-testing readings of bacteria w.r.t. AMR drugs.

The source instance contains the data shown in Table 1 and Table 2. The latter table shows two relations containing information about hurricanes and farms that could be relevant to the data in Table 1.

HurricaneInfo	Name	Location	Date
	Harvey	Houston	[8/25/17, 8/29/17]

FarmInfo	Farm Name	Location
	Louisburg Farm	Louisburg
	Houston Farm	Houston

Table 2: Assumed information of hurricane and farms.

2.1.2 Target Schema and Chasing S-T Tgds

The target instance uses the schema shown in Table 3. The intent is to describe the *TET MIC* readings that were collected at the locations of, and during the occurrence of, hurricanes.

HurricaneReadings					
Date	Hurricane	Loc	Isolate	Serotype	TET MIC

Table 3: Schema for readings of isolates that are co-located and co-occurring with hurricanes.

To exchange the source instance stored in the relations described by Table 1 and Table 2 into the target schema shown in Table 3, the following temporal

s-t tgds can be used.

$$\begin{aligned}
 & \text{AnimalReadings}(t_1, farm, sid, stype, mic) \\
 & \wedge \text{HurricaneInfo}(wname, loc, t_2) \\
 & \wedge \text{FarmInfo}(farm, loc) \\
 & \wedge t_1 \subseteq t_2 \\
 & \rightarrow \text{HurricaneTestReadings}(t_1, wname, loc, sid, stype, mic)
 \end{aligned} \tag{1}$$

This temporal s-t tgds serves as a filter that only picks up the *TET MIC* readings that are directly related to hurricanes. On the left-hand side, the tgds uses multiple temporal variables (t_1, t_2) and Allen’s interval relation $[1] \subseteq$ (*starts or during or finishes*) to compare them. Source values, e.g., *wname*, would be “exchanged” to the target instance only if the temporal comparison evaluates to **True**.

Chasing the source with the s-t tgds of Eq.(1) results in the target instance shown in Table 4.

HurricaneTestReadings					
Date	Hurricane	Loc	Isolate	Serotype	TET MIC
[8/27/17, 8/27/17]	Harvey	Houston	S2020	S.Anatum	>32

Table 4: The target instance resulting from chasing the source with the s-t tgds of Eq.(1).

2.1.3 Temporal Query Answering

AMR experts can use the relation with the schema shown in Table 3 to retrieve the resistance-test readings that are directly related to hurricanes during a certain year, e.g., 2017. This can be done by posing Query 2.1.3 on the instance shown in Table ??.

Query 2.1.3 SELECT Hurricane, Serotype, TET.MIC, ‘2017’ AS YEAR
 FROM HurricaneTestReadings
 WHERE Date \subseteq [1/1/2017, 12/31/2017]

The result of evaluating this query on the data of Table 4 is shown in Table 5.

Hurricane	Serotype	TET MIC	YEAR
Harvey	S.Anatum	>32	2017

Table 5: Answers to Query 2.1.3 w.r.t. the data shown in Table 4.

2.2 Example 2

2.2.1 Source Schema and Instance

This example uses data from two AMR spreadsheets of the same schema, see Table 6 (with information on resistance-testing readings collected on animals) and Table 7 (with information on readings collected on humans). The attribute *TET RSI* interprets the value of *TET MIC* into three categories: *Resistant (R)*, *Susceptible (S)*, and *Intermediate (I)*.

AnimalReadings				
Date	Sal ID	Serotype	TET MIC	TET RSI
[10/22/09, 10/22/09]	S481	S.Typhimurium	>32	R
[2/22/10, 2/22/10]	S939	S.Newport	<4	S

Table 6: Resistance-testing readings of bacteria to drug *TET* on animals.

HumanReadings				
Date	Sal ID	Serotype	TET MIC	TET RSI
[7/8/09, 7/8/09]	HS6	S.Newport	>32	R
[10/22/09, 10/22/09]	HS302	S.Typhimurium	>32	R

Table 7: Resistance-testing readings of bacteria to drug *TET* on humans.

2.2.2 The Target Schema and S-T Tgds

The intent of the design of the target schema is to exhibit information on the kinds of bacteria (*Serotype*) that have the same level of resistance to the *TET* drug between animals and humans during the same time periods. The schema is shown in Table 8.

PrevalentResistance	Serotype	Resistance	Date

Table 8: Bacteria that are resistant in both animals and humans.

To exchange the source instance into the target schema, we use the following temporal s-t tgds.

$$\begin{aligned}
 & \text{AnimalReadings}(t_1, sid_1, stype, mic_1, r) \\
 & \wedge \text{HumanReadings}(t_2, sid_2, stype, mic_2, r) \\
 & \wedge t_1 \text{ overlaps } t_2 \\
 & \rightarrow \text{PrevalentResistance}(stype, r, t_1 \wedge t_2)
 \end{aligned} \tag{2}$$

This temporal s-t tgds performs as a filter to help users retrieve the kinds of bacteria that are at the same resistance level (i.e., both resistant / susceptible /

intermediate) in both animals and humans at the same time. This information is of importance in the AMR domain, as prevalent resistance phenomena can highlight potential abuses of AMR drugs among animals and humans, and thus can help domain experts make mitigation decisions. On the left-hand side, the tgd of Eq.(2) uses Allen’s interval relation *overlaps* to compare two temporal variables, t_1 and t_2 . Individual values of the variables will be transmitted to the target instance only if the temporal comparison evaluates to **True**.

Chasing the source data shown in Tables 6 and 7 with the s-t tgd of Eq.(2) results in the target instance shown in Table 9.

PrevalentResistance		
Serotype	Resistance	Date
S.Typhimurium	R	[10/22/09, 10/22/09]

Table 9: The target instance resulting from the chase of the data of Tables 6–7 with the s-t tgd of Eq.(2).

2.2.3 Temporal Query Answering

AMR experts can use the relation of Table 9 to retrieve the serotypes of bacteria that showed resistance in both animals and humans during a certain year, e.g. 2009. This can be done by posing Query 2.2.3 on the instance shown in Table 9.

Query 2.2.3 SELECT Serotype, ‘2009’ AS YEAR, ‘TET’ AS DRUG
 FROM PrevalentResistance
 WHERE Date \subseteq [1/1/2009, 12/31/2009]

The answers to the query are shown in Table 10.

Serotype	YEAR	DRUG
S.Anatum	2009	TET

Table 10: Answers to Query 2.2.3 on the data shown in Table 9.

2.3 Example 3

2.3.1 Source Schema and Instance

Examples 1–2 showed how temporal s-t tgds could arise in exchanging raw AMR data stored in spreadsheets. We now show an example of s-t tgds arising in exchanging summarized AMR information.

In this example, the source instance, shown in Table 11, summarizes AMR information by year. The attribute *# isolates* indicates the number of isolates that were tested for resistance for a kind of bacteria (*Serotype*) within the year

specified by *Year*. *% resistance* represents the percentage of the tested isolates showing resistance to the AMR drug specified by *AMR Drug*. The data in this table are generated by aggregating raw spreadsheet data and combining the result with summary data published annually in governmental reports such as DANMAP reports [9].

AnimalResistance				
Year	Serotype	# isolates	AMR Drug	% resistance
[2009, 2009]	S.Newport	333	AMP	1.5
[2015, 2015]	S.Newport	232	AMP	5.6

Table 11: Yearly AMR information on animals.

2.3.2 The Target Schema and S-T Tgds

The intent of the target instance is to store information on the kinds of bacteria (*Serotype*) that are increasingly resistant to *AMR* drugs in animals. It uses the schema shown in Table 12.

IncreasingResistance	Serotype	AMR Drug	Year

Table 12: Bacteria that are increasingly AMR resistant in animals: the schema.

The following temporal s-t tgds is used to exchange the source instance into the target schema.

$$\begin{aligned}
& AnimalResistance(t_1, stype, iso_1, drug, perc_1) \\
& \wedge AnimalResistance(t_2, stype, iso_2, drug, perc_2) \\
& \wedge t_1 < t_2 \\
& \wedge perc_1 < perc_2 \\
& \rightarrow IncreasingResistance(stype, drug, t_1 \vee t_2)
\end{aligned} \tag{3}$$

This temporal s-t tgds performs as an aggregator that summarizes the resistance information for serotypes of bacteria, while retaining only those serotypes that show increasing resistance to drugs. The target instance can aid users in determining the potentially excessive uses of drugs. On the left-hand side, the tgds of Eq.(3) uses Allen’s interval relation $<$ (*before*) to compare temporal variables t_1 and t_2 . $stype$ will be exchanged to the target instance only if the temporal comparison evaluates to **True**.

Chasing the source data with the s-t tgds of Eq.(3) results in the target instance shown in Table 13.

IncreasingResistance	Serotype	AMR Drug	Year
	S.Newport	AMP	[2009, 2015]

Table 13: The target instance resulting from the chase of the data of Table 11 with the s-t tgd of Eq.(3).

2.3.3 Temporal Query Answering

AMR experts can use the relation of Table 13 to retrieve the serotypes of bacteria for which resistance increased between the year of 2012 to 2019. This can be done by posing Query 2.3.3 on the instance shown in Table 13:

Query 2.3.3 SELECT Serotype, AMR Drug, Year
 FROM IncreasingResistance
 WHERE Year *overlaps* [2012, 2019]

The answers to the query are shown in Table 14.

Serotype	AMR Drug	Year
S.Newport	AMP	[2012, 2015]

Table 14: Answers to Query 2.3.3 on the data shown in Table 13.

2.4 Example 4

In this example, we illustrate how temporal GLAV mappings arising in temporal exchange of AMR data.

2.4.1 Source Schema and Instance

The source instance used in this example is shown in Table 15; the attribute *Date* in the relation *SampleInfo* indicates the collection date for a sample, and attribute *Date* in the relation *AMRDrugs* represents the date at which an AMR drug became available to the local laboratory performing resistance testing.

SampleInfo	Sample	Location	Date
	PE9	Louisburg Farm	[5/19/09, 5/19/09]

AMRDrugs	Name	Class	Date
	Ampicillin	Penicillins	[1/1/09, 1/1/09]
Tetracycline	Tetracyclines	[8/8/09, 8/8/09]	

Table 15: The source instance in Example 4.

2.4.2 The Target Schema and S-T Tgds

The intent of the target instance is to store resistance-testing readings for the bacteria present in the collected samples, with respect to all the AMR drugs that had become available to the laboratory before the collection dates for the samples. The schemata of the target relations are shown in Table 16.

	IsolateInfo		Sample	
	Isolate	AMR Drug	Reading	Date
Readings				

Table 16: Resistance-testing readings of collected samples: the schema.

To exchange the source instance into the target schema, we use the following temporal s-t tgd.

$$\begin{aligned}
 & SampleInfo(sp, loc, t_1) \\
 & \wedge AMRDrugs(name, cls, t_2) \\
 & \wedge t_2 \leq t_1 \\
 & \rightarrow \exists iso \exists r (IsolateInfo(iso, sp) \wedge Readings(iso, name, r, t_1))
 \end{aligned} \tag{4}$$

This temporal s-t tgd performs as a value generator, specifying the knowledge that for the AMR drugs that were already available to the laboratory at the collection time for a sample, the laborstory would perform resistance testing on the sample with respect to each of the drugs. On the left-hand side, the tgd uses Allen’s interval relation \leq (*before or equal*) to compare temporal variables t_1 and t_2 . When the temporal comparison evaluates to **True**, the values of sp , $name$, and t_1 will be “exchanged” into the target instance, and labeled nulls would be generated for the existentially quantified variables.

Chasing the source data with the s-t tgd of Eq.(4) results in the target instance shown in Table 17.

	IsolateInfo		Sample	
	Isolate	AMR Drugs	Readings	Date
Readings	I_1	Ampicillin	R_1	[5/19/09, 5/19/09]

Table 17: Target instance after chasing the source data with the s-t tgd of Eq.(4).

2.4.3 Temporal Query Answering

AMR experts can use the relation with the schema shown in Table 16 to retrieve information, including testing samples and dates, on the AMR drugs that have been used in resistance testing. This can be done by posing Query 2.4.3 on the instance shown in Table 17.

Query 2.4.3 `SELECT AMR Drugs, Sample, Date
FROM IsolateInfo, Readings`

The answers to the query are shown in Table 18.

AMR Drug	Sample	Date
Ampicillin	PE9	[5/19/09, 5/19/09]

Table 18: Query answers to Query 2.4.3.

2.5 Example 5

In this example, we provide another illustration of temporal GLAV mappings arising in temporal exchange of AMR data.

2.5.1 Source Schema and Instance

The source instance used in this example is shown in Table 19; the *Dates* attribute in the *DrugUsage* relation represents a period of consecutive days of administering an AMR drug to an animal, and the *Dates* attribute in the *SampleCollection* relation indicates a period of consecutive days for collecting samples from an animal. Note that the process of sample collection can span several days. In such cases, one sample would be collected per day for each day in the given time period, and the samples collected during this period will be labeled as an individual-sample collection. (Cf. the setting that was used in [10], as illustrated in their Table S1.)

DrugUsage	Drug Name	Animal	Dates
	Ampicillin	P1	[5/10/09, 5/24/09]
SampleCollection	Collection	Animal	Dates
	C1	P1	[5/19/09, 5/26/09]

Table 19: The source instance of Example 5.

2.5.2 Target Schema and S-T Tgds

The intent of the target instance is to store information on how bacteria in the bodies of animals react to the AMR drugs used on the animals. If a drug was used on an animal and a sample was collected on the animal on the same day, then isolates of the bacteria would be detected in the sample on that day, a bacteria-resistance test would be performed on the detected isolates on the same day, and, finally, a resistance-testing reading would be generated on the same day to indicate the reactions of the bacteria to the drugs.

The target schema contains the two relation schemas shown in Table 20. The *Isolate* attribute in each relation represents the isolates (cultures) of bacteria

detected from a sample in a sample collection. The *Detecting Dates* attribute represents the dates of detection of isolates, and the *Testing Dates* attribute denotes the dates of performing resistance tests and generating the test results.

DetectedIsolates			
Isolate	Sample Collection	Detecting Dates	

PostDrugUseReadings			
Isolate	AMR Drug	Reading	Testing Dates

Table 20: Post-drug-usage resistance-testing readings from samples: the schema.

To exchange the source instance into the target schema, we use the following temporal s-t tgdt.

$$\begin{aligned}
 & DrugUsage(d, a, t_1) \wedge SampleCollection(s, a, t_2) \wedge t_1 \text{ overlaps } t_2 \\
 & \rightarrow \exists i \exists r (DetectedIsolate(i, s, t_1 \wedge t_2) \\
 & \wedge PostDrugUseReadings(i, d, r, t_1 \wedge t_2))
 \end{aligned} \tag{5}$$

This temporal s-t tgdt performs as a value generator, specifying the knowledge that there must be a resistance-testing reading generated for the day on which an AMR drug was used and a sample was collected, and that readings generated on different days can be different. On the left-hand side, the tgdt uses Allen’s interval relation *overlaps* to compare temporal variables t_1 and t_2 . When the temporal comparison evaluates to **True**, the values of the variables b and s will be “exchanged” into the target instance, and labeled nulls would be generated for the existentially quantified variables.

Chasing the source instance with the s-t tgdt of Eq.(5) results in the target instance shown in Table 21.

DetectedIsolates			
Isolate	Sample Collection	Detecting Dates	
$I_1^{[5/19/09,5/24/09]}$	C1	[5/19/09, 5/24/09]	

PostDrugUseReadings			
Isolate	AMR Drugs	Readings	Testing Dates
$I_1^{[5/19/09,5/24/09]}$	Ampicillin	$R_1^{[5/19/09,5/24/09]}$	[5/19/09, 5/24/09]

Table 21: Target instance after chasing the source data with the s-t tgdt of Eq.(5).

2.5.3 Temporal Query Answering

AMR experts can use the relation with the schema shown in Table 20 to retrieve the readings that have been collected on resistance testing, together with

their testing samples, drugs, and dates, during a specific time period, e.g., from 5/1/2009 to 5/31/2009. This can be done by posing Query 2.5.3 on the instance shown in Table 21.

Query 2.5.3 SELECT Readings, Sample Collection, AMR Drugs,
Testing Dates
FROM DetectedIsolates, PostDrugUseReadings
WHERE Testing Dates \subseteq [5/1/09, 5/31/09]

The answers to the query are shown in Table 22.

Readings	Sample Collection	AMR Drugs	Testing Dates
$R_1^{[5/19/09, 5/24/09]}$	C1	<i>Ampicillin</i>	[5/19/09, 5/24/09]

Table 22: Answers to Query 2.5.3 on the data shown in Table 21.

2.6 Example 6

In this example, we provide another illustration of temporal GLAV mappings arising in temporal exchange of AMR data. This example exhibits existentially quantified temporal variables on the right-hand side of s-t tgds, as well as multidimensional time and Allen’s interval comparisons on the right-hand side. This example is closely connected to the previous example, as both Example 5 and Example 6 share the same source instance and are demonstrating the same resistance-testing scenario.

2.6.1 Source Schema and Instance

See the source schema and instance for Example 5 (Section 2.5).

2.6.2 Target Schema and S-T Tgds

The target schema here differs from that of Example 5, where we assumed that resistance-testing readings are generated on the same day on which an AMR drug is used and a sample is collected. However, real-life AMR resistance testing usually requires overnight incubation [2], which means that test results would be available on the next day after the date of the resistance testing.

To reflect in the data information about this scenario, as shown in Table 23, in this example we use two temporal attributes in the target relation *PostDrugUseReadings*, which captures temporal information about resistance testing. Here, the *Starting Date* attribute represents the date on which resistance testing starts (this is the date of using a drug and collecting a sample), and the *Date of Posting Results* attribute denotes the date on which test results are available. (It is logical to assume that the value of *Date of Posting Results* is greater than the value of *Starting Date*.)

DetectedIsolates				
Isolate	Sample Collection	Detecting Dates		

PostDrugUseReadings				
Isolate	AMR Drug	Reading	Starting Date	Date of Posting Results

Table 23: Post-drug-usage resistance-testing readings of collected samples, with multiple temporal attributes: the schema.

To exchange the source instance into the target schema, we use the following temporal s-t tgd.

$$\begin{aligned}
& DrugUsage(d, a, t_1) \wedge SampleCollection(s, a, t_2) \wedge t_1 \text{ overlaps } t_2 \\
& \rightarrow \exists i \exists r \exists t_3 (DetectedIsolates(i, s, t_1 \wedge t_2) \\
& \wedge PostDrugUseReadings(i, d, r, t_1 \wedge t_2, t_3) \\
& \wedge (t_1 \wedge t_2) \text{ overlaps } t_3)
\end{aligned} \tag{6}$$

This temporal s-t tgd performs as a value generator, specifying the knowledge that: (i) there must be a resistance test started on the same day as an AMR drug was used and a sample was collected; that (ii) there would be a delay between the dates of starting tests and the dates of posting results; and that (iii) readings generated on different days can be different.

On the left-hand side, the tgd uses Allen’s interval relation *overlaps* to compare temporal variables t_1 and t_2 . When the temporal comparison evaluates to **True**, the values of the variables b and s will be “exchanged” into the target instance, and labeled nulls would be generated for the existentially quantified variables, for the time interval that witnesses t_1 , t_2 and t_3 .

Chasing the source instance with the s-t tgd of Eq.(6) results in the target instance shown in Table 24.

DetectedIsolates		
Isolate	Sample Collection	Detecting Dates
$I_1^{[5/19/09,5/24/09],[5/20/09,5/25/09]}$	$C1$	$[5/19/09, 5/24/09]$

PostDrugUseReadings		
Isolate	AMR Drugs	Readings
$I_1^{[5/19/09,5/24/09],[5/20/09,5/25/09]}$	$Ampicillin$	$R_1^{[5/19/09,5/24/09],[5/20/09,5/25/09]}$
	Starting Dates	Date of Posting Results
	$[5/19/09, 5/24/09]$	$[5/20/09, 5/25/09]$

Table 24: Target instance after chasing the source data with the s-t tgd of Eq.(6).

2.6.3 Temporal Query Answering

AMR experts can use the relation with the schema shown in Table 23 to retrieve the readings that have been collected on resistance testing, with start and end dates during a specific time period, e.g., from 5/1/2009 to 5/31/2009, along with the associated test samples, drugs, and related dates. This can be done by posing Query 2.6.3 on the instance shown in Table 24.

Query 2.6.3 SELECT Readings, Sample Collection, AMR Drugs,
Starting Dates, Date of Posting Results
FROM DetectedIsolates, PostDrugUseReadings
WHERE Starting Dates \subseteq [5/1/09, 5/31/09] AND
Date of Posting Results \subseteq [5/1/09, 5/31/09]

The answers to this query on the data of Table 24 are shown in Table 25.

Readings	Sample Collection	AMR Drugs
$R_1^{[5/19/09,5/24/09],[5/20/09,5/25/09]}$	$C1$	<i>Ampicillin</i>
Starting Dates	Date of Posting Results	
[5/19/09, 5/24/09]	[5/20/09, 5/25/09]	

Table 25: Answers to Query 2.6.3 on the data shown in Table 24.

References

- [1] James F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.
- [2] Jennifer M Andrews. Determination of minimum inhibitory concentrations. *Journal of antimicrobial Chemotherapy*, 48(suppl_1):5–16, 2001.
- [3] Jan Chomicki and David Toman. Temporal databases. In *Handbook of Temporal Reasoning in Artificial Intelligence*, pages 429–467. 2005.
- [4] R. Fagin, P. G. Kolaitis, R. J. Miller, and L. Popa. Data Exchange: Semantics and Query Answering. *Theoretical Computer Science*, 336(1):89–124, 2005.
- [5] Ronald Fagin, Phokion G Kolaitis, Renée J Miller, and Lucian Popa. Data exchange: semantics and query answering. *Theoretical Computer Science*, 336(1):89–124, 2005.
- [6] Claudio Gutierrez, Carlos A Hurtado, and Alejandro Vaisman. Introducing time into rdf. *IEEE Transactions on Knowledge and Data Engineering*, 19(2):207–218, 2007.

- [7] Pei-Yu Hou, Jing Ao, and Rada Chirkova. Building a data-analytics workflow for antimicrobial-resistance studies: An experience report, 2019. NCSU technical report TR-2019-7; available online at <https://www.csc.ncsu.edu/research/tech/reports.php>.
- [8] Nuno Lopes, Axel Polleres, Umberto Straccia, and Antoine Zimmermann. Anql: Sparqling up annotated rdfs. In *International Semantic Web Conference*, pages 518–533. Springer, 2010.
- [9] Danish Integrated Antimicrobial Resistance Monitoring. Research programme (DANMAP). Use of antimicrobial agents and occurrence of antimicrobial resistance in bacteria from food animals, food, and humans in Denmark. <https://danmap.org>.
- [10] Pingfeng Yu, Avery Zaleski, Qilin Li, Ya He, Kris Mapili, Amy Pruden, Pedro JJ Alvarez, and Lauren B Stadler. Elevated levels of pathogenic indicator bacteria and antibiotic resistance genes after hurricane harvey’s flooding in houston (supporting information). *Environmental science & technology letters*, 5(8):481–486, 2018.