On Modeling Political Systems to Support the Trust Process

Carlos Laufer[1[+552135271500]] and Daniel Schwabe[1 [+552135271500]]

[1] Dept. of Informatics, PUC-Rio
R. M. de S. Vicente, 225, Rio de Janeiro, RJ 22453-900, Brazil
laufer@globo.com, dschwabe@inf.puc-rio.br

Abstract. Modern political life is heavily dependent and influenced by claims about facts made about participants in a Political System (seen here as sets of agents and organizations that govern a society). Such claims are commonly found in news stories, tweets, and social media postings, and, depending on each user, they may be considered "true" or "false". The decision to act based on a fact contained in a claim is dependent on the agent's trust on the claim, i.e., that the fact being stated is indeed "true". In this paper, we propose a framework to describe the trust *process*, which can serve as a a support for various possible models of content trust. Within this framework, we present POLARE, an ontology to describe a Political System, which includes provenance information, and illustrate how it may provide information to feed the trust process in online news stories.

As a case study, we illustrate the approach using a Knowledge Graph created with such claims, which can be queried in order to better understand the existing relations between agents in the Political System in Brazil. In order to use claims in this Knowledge Graph, users will typically require provenance information to make trust judgements about them.

Keywords: trust, provenance, linked data, political systems, news.

1 Introduction

Modern society has become heavily influenced by information (and misinformation) that flows in news sites and social networks in the Internet [10]. There are many studies carried out in several disciplines attempting to characterize and understand the spread of information in the cybersphere, and how this affects society[1] [3,27,26]. An arguably more recent phenomenon has been the spread of "fake news", actually a term used to refer to several different misuses of information, as postulated by Wardle in [29]. This has also been the focus of much research and many initiatives [9,1,18,6], including manual fact-checking by sites like Snopes[2], Politifact[3],

[1] See for example, https://www.pheme.eu/ and https://revealproject.eu.
[2] http://www.snopes.com
[3] http://www.politifact.com/

FullFact[4], FactCheck[5], and Crosscheck[6], and automated approaches such as Factmata[7], Hoaxy[25] and others [7,28,22].

This phenomenon highlights an often-implicit assumption in most information systems, including websites – the fact that the data it uses is assumed to be "true", in the sense that the application can rely on this data to perform whatever computation it needs to do in order to fulfill its purposes. The advent of user-contributed data raised the issue of assessing its "data quality" [19], and the fact that data, in reality, expresses a belief or opinion of some agent. This becomes mostly evident in case of online reviews online and social media, often with direct effect on commercial success (e.g. [14]).

The original vision for the Semantic Web already envisioned a "Trust" layer, although its emphasis was more on authentication and validation, and static trust measures for data. There have been many efforts in representing trust, including computational models; a general survey can be found in [22], [5] presents an excellent earlier survey for the Semantic Web, and [24] surveys trust in social networks. The advent of Linked Data highlighted the vision that facts in Semantic Web should be regarded as claims rather than hard facts [7].

In most if not all research surveyed in [22], there is an implicit model of what is the trust process that is being adopted or assumed, which is never spelled out precisely. In this paper, we propose a framework to describe the trust *process*, which can incorporate various alternative models of content trust (as proposed in [11]). Within this framework, we present POLARE, an ontology to describe a Political System, which includes provenance information, and illustrate how it may provide information to feed the trust process in online news stories.

The remainder of this organized as follows. Section 2 presents a model of the trust process; Section 3 presents POLARE; Section 4 shows an example of its use, and Section 5 discusses related work, and Section 5 draws some conclusions and points to future work.

2 A Framework for the Trust Process

In this section, we characterize the trust process that underlies the use or consumption of data/information on the web. Figure 2 shows a diagram of our proposed framework for the trust process model.

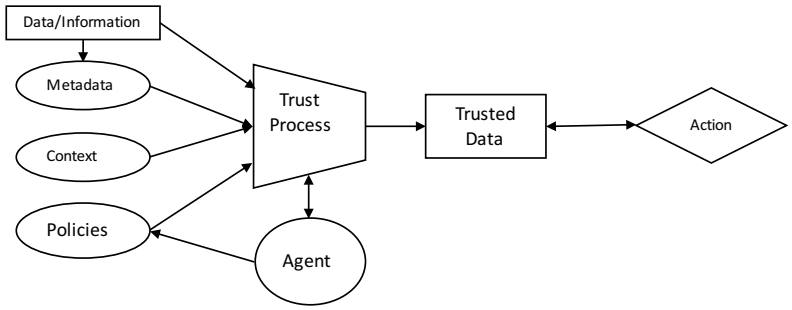

Figure 1 – A Framework of the Trust Process

We focus here on the cases where an *Agent* needs to act, i.e., do some computation, make a decision, or take some *Action*. The agent must act based on some *Data/Information* items which, clearly, it must trust – the *Trusted Data*. Contrary to how it used to be in pre-Internet times, the

[4] https://fullfact.org/
[5] https://crosscheck.firstdraftnews.com/france-en/
[6] http://factcheck.org/
[7] http://factmata.com/

Data/Information items to be used by the agent may come from several sources, and it is not always clear (to the *Agent*) what is the quality of this *Data*, or the trustworthiness of the *Information* it containes. Therefore, the *Agent* must apply a *Trust Process* to filter the incoming *Data/Information* items and extract the *Trusted Data* items to be used by the *Action*. From this point of view, an item is considered as the smallest indivisible element that can be used in the *Trust Process,* and may have an internal structure when used by the *Action*.

This *Trust Process* can be based on a multitude of different signals, some of which we have singled out in the diagram in Figure 1, namely, the *Metadata* which describes various properties of the *Data/Information* items, and the *Context* in which the *Action* will take place. The criteria used in the *Trust Process* are expressed by *Policies* determined by the *Agent*. Notice that in this framework, the *Context* contains any arbitrary information items used by the *Policies*, in addition to *Metadata* and the *Data/Information* items themselves.

Many, if not most, of the trust models discussed in the surveys mentioned in Section 1 [5,19,24] can be regarded as providing models or representations for one or more of the elements of this trust framework, such as different representations for the metadata, or specification language for the policies, or particular types of context information that can be added to the data. Regardless of these models, it should be clear that, as far as the *Action* is concerned, the whole process is *binary*: An incoming *Data/Information* item is either accepted and inserted into the *Trusted Data*, or it is not – there are no "half filters". In other words, when time comes to *Act*, either the *Agent* uses that *Data/Information* item, or it does not, it can't "use it partially" – an element is considered indivisible from the point of view of the *Trust Process* [4].

Thus, the actual trusting process should not be confused with models which may attribute non-discrete or continuous values to the trustworthiness assigned to data/information items that are used to determine if the incoming data/information items should be filtered or not.

Another point to notice is that, in this framework, only the *Agent* determines the policies it wishes to apply to incoming *Data/Information*. The metadata associated with the incoming Data/Information may contain *Data/Information* about the publisher or provider of a *Data/Information* item, and the *Policies* may take this into account in the filtering process.

In contrast, privacy issues would require adding a similar set of metadata and policies to the *published* Data/Information that will be consumed by the *Agent*, which would act as an additional filter, applied before the Data/Information is made available to the *Agent*. The privacy *Policies* of a Publisher may use information about the (requesting) *Agent*, the *Action* and the *Context*.

In this work, we will focus on specific types of *Data/Information* – those pertaining to Political Systems, and one specific type of metadata, *Provenance*, in the context of news media. In the next section, we discuss the trust process in this context, and propose a model for information and metadata to support it.

3 Modeling Political Agents, Relationships and Provenance

The nature of information in news and social media must, per force, be seen as containing claims made by someone about some "fact" (or "statement"), i.e., someone (an agent, i.e. a person or organization) claims a certain fact to be true. It is up to the reader to accept the truth of these claimed facts.

In order to accept the truth of a claim, the reader either already knows (i.e. believes in) the truth of the claimed fact, or s/he must somehow establish trust in the agent making the claim.

Whereas there are many definitions of trust [16,17], we follow the view that trust is "the expectation that some agent will act in a predictable, expected way, in a given situation". Therefore, to accept a certain claim about a hitherto unknown fact entails either establishing the truth through direct observation (or a chain of interdependent observations), or trusting the source of the claim, possibly based on some evidence. The trust on the source of the claim then implies

that one expects the claim to be true because that source has been consistently (predictably and expectedly) correct in the past.

If, however, the source is not trusted, one may require some evidence to serve as the basis to establish the truth of the claim. Such evidence is, in turn, another claim made by some third party, which the claimant believes is more trusted by the first user than her/himself. In this fashion, a chain of interdependent claims is established, such that the truth of the original claim relies on the trust on the agent at the end of the chain. In modern societies, certain public agents (e.g., Vital Record Office, Notary Public) have, by convention and mutual agreement among its members, the so-called "public faith", meaning that whatever they claim (about certain specific facts) is deemed "true". For instance, John's mother claims he was born in Brazil, but for many purposes, this claim is only accepted as true if she can produce a birth certificate issued by the proper authority stating that he was indeed born in Brazil.

From the discussion above, it is safe to say that provenance data about information published in the cybersphere is crucial to enable agents to consume this information – i.e., accept it as true and eventually act based on this. This has been clearly recognized by many researchers and scholars studying the "fake news" phenomenon [13,15].

In [5] Wardle proposes 7 types of mis- and disinformation:

1. Satire or parody – when there is no intention to cause harm, but can be misinterpreted and cause harm;
2. Misleading content – Use of information to frame an issue or individual in a misleading way;
3. Imposter content – When a genuine source is impersonated;
4. Fabricated content – New content that is entirely false, designed to deceive and do harm;
5. False connection – Use of headlines, visuals or captions that don't support the actual content;
6. False context – Sharing of genuine content with false contextual information;
7. Manipulated content – When genuine information or imagery is manipulated to deceive.

A large part of fact checking, especially when automated, focuses on claims that are in clear contradiction with verifiable facts – type 4 above.

We have been involved with an initiative to publish a database about Political Agents in Brazil in the form of Linked Data, named "Se Liga na Politica"[8] (SLNP). The data in this database is obtained from several sources, in both automated and non-automated ways. Most of the automated extraction is made from official sources, such as the open data published by the House of Deputies and by the Senate. In addition to such sources, data may also be contributed by individuals, in wiki fashion. One of the main usages for this database is to provide context information for news stories, to allow readers to establish trust in the claimed facts based on their own criteria. We present here the rationale and design decisions for the underlying ontologies used in this database, starting first with our motivations.

We regard a Political System as sets of agents and organizations that govern a society. Our goal in the SNLP project is to help making the existing relationships between political agents in a Political System explicit and thus subject to analysis. We do so by providing a Linked Data database where such relationships appear as data items. Given the multiplicity of sources, and the nature of the subject matter, this database is designed so that facts are seen as claims made by some agent, and therefore provenance information becomes a "first class citizen" of the domain.

We first describe the "domain" ontology dealing with agents and their relationships, and then describe how provenance is integrated.

---

[8] The expression "Se liga" in Portuguese has a colloquial meaning of "be aware", "pay attention to", as well as "connect yourself". In Portuguese, it reads both as "pay attention to Politics" and "Connect yourself to Politics". A third (indirect) meaning is the reference to Linked Data.

3.1 POLARE – A Political Agents Relationship Ontology

The central concepts in POLARE are People, Organizations and the Posts that people occupy in Organizations. The main terms of the POLARE vocabulary come from W3C's ORG Ontology[9] and are shown in Figure 2. It is important to notice the reification of the relation "occupies" between Person and Post via the Membership class, which is fundamental to allow us to represent properties of this relation, as will be discussed later.

We have added the "polare:hasPost" relation linking "org:Membership" with "org:Post" because we need to represent that a Post might be occupied by different persons in different periods of time.[10]

Both Post and Membership may have start and end dates associated to them. The dates associated to a Post refer to a time period when the Post exists, for example, for a post in the House of Representatives, it corresponds to a Legislature, which defines the mandate of the elected person. The Membership dates refer to the period in which the Person actually occupies the Post, as it is possible for a Person to temporarily leave the Post for a short period of time within the mandate.

For example, every four years there is a presidential election in Brazil to select a President and a Vice-President to lead the country during the succeeding four years. Therefore, there is a President Post (P1) and a Vice-President Post (P2) in the Federal Government (an Organization) that have established start and end dates. We have to create a Membership (M1) for the person that was elected for President with the same start date as P1, and a Membership (M2) for the person that was elected for Vice-President with the same start date as P2.

In some situations, for example, when the President has to travel outside the country, the Vice-President assumes the Post of President for the period of time that the elected President is abroad. In that case, we have to add the end date to M1 with the initial date of the travel. We have to create a new Membership (M3), related to P1, with a start date equal to the initial date of the travel, and assign it to the person occupying the Vice-President post P2. When the President returns, we have to add the end date of M3 with the final date of the travel. In addition, we have to create a new Membership (M4), related to P1, with a start date equal to the final date of the travel, and allocate it to the President again. Thus, after the travel period, the database will have three Memberships (M1, M3, M4) related to the same Post (P1).

[9] https://www.w3.org/TR/vocab-org/

[10] Notice that "Post" here refers to a "Position", not to a blog post.

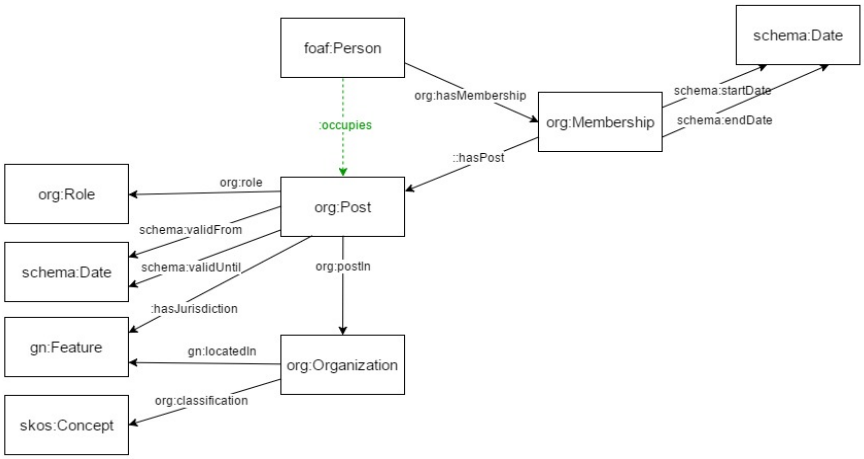

Figure 2- The main concepts in the POLARE ontology.
Inferred relations are shown as dotted lines.

Several indirect relationships of interest between Persons can be captured simply by the fact that these Persons occupy Posts in the same Organization. For example, the fact that two congressmen belong to the same party (which is an Organization), or that they were once colleagues in a company in the private sector or in a department in the executive branch.

Another important kind of relationships between persons are direct kinships, which are modeled in POLARE as Direct Relationships, shown in Figure 3. Once more, since it is necessary to represent some properties of these relationships, such as the dates involving a marriage, we use the traditional class approach for reifying the relationship property. The directRelProp property allows specifying what is the kinship relation, typically using the Rel vocabulary[11].

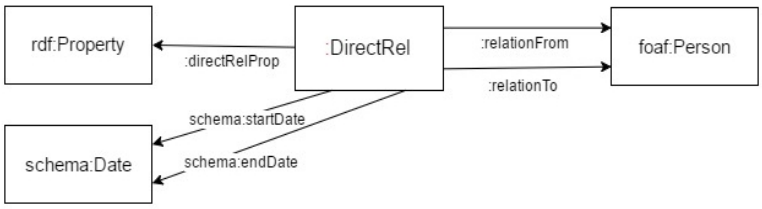

Figure 3 – Direct relationships between Persons in POLARE

In Political Systems, and also in organizations, it is important to know if some Post was filled through a Referral, i.e., that some Agent indicated (nominated) a Person to occupy a Post. POLARE can represent which Agent has referred some Person to occupy a Post in an Organization, shown in Figure 4. We also reified the "refers" property using a Referral class, to allow including data related to the Referral.

[11] http://vocab.org/relationship/

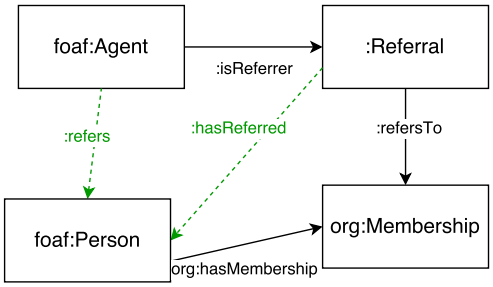

Figure 4 – Modeling Referrals in POLARE

3.2 Heart - a Claim Provenance ontology pattern

As already discussed, a critical information to establish trust is data related to the provenance of that information, since acceptance of the truth of a claim relies on the trust one has on the Agent making the claim. For the Political Systems domain, the focus is on the relationships between agents, which are modeled in POLARE using the Membership, Direct Relation and Referral classes.

The PROV-O Ontology[12] provides a vocabulary to represent provenance information, which can be used in a variety of ways to represent diverse aspects of provenance. For our purposes, especially because we allow individual users to enter information (including provenance) in the database, it is important to ensure that the provenance information is recorded in a uniform way, providing enough information to allow the trust process. To achieve this, we resort to the notion of a pattern, similar in spirit to the PROV template proposed in [20]. The advantage of using PROV Patterns is to extend the ontology with a well-formed structure using the PROV vocabulary, allowing the construction of queries that use PROV data mixed with the data of the main ontology.

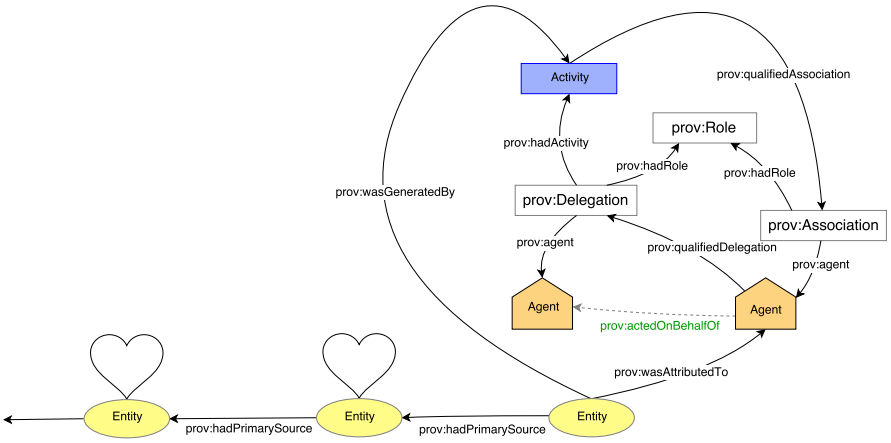

Figure 5 – The PROVHeart Provenance Pattern. A provenance chain can be established using the prov:hasPrimarySource property. The "heart" symbol represents an instance of this pattern.

We define a provenance pattern called PROVHeart, shown in Figure 5. PROVHeart is a simple pattern that models the provenance of an Entity that was generated by an Activity attributed to an Agent with a Role in that Activity, using another Entity as a source, and acting on

[12] https://www.w3.org/TR/prov-o/

behalf of another Agent. Provenance chains can be built using the prov:hadPrimarySource relation between Entities that serve as the "justification" or "basis" for each claim, forming the basis for the trust chain discussed earlier. Notice also that "prov:actedOnBehalfOf" is represented as a property chain property, because this relation has to be linked to the Activity that produced the particular Entity the provenance refers to. In other words, "prov:actedOnBehalfOf" is seen as applying to specific activities, and not as a "blanket statement" that *always* holds for all Activities carried out by an Agent, relative to another Agent.

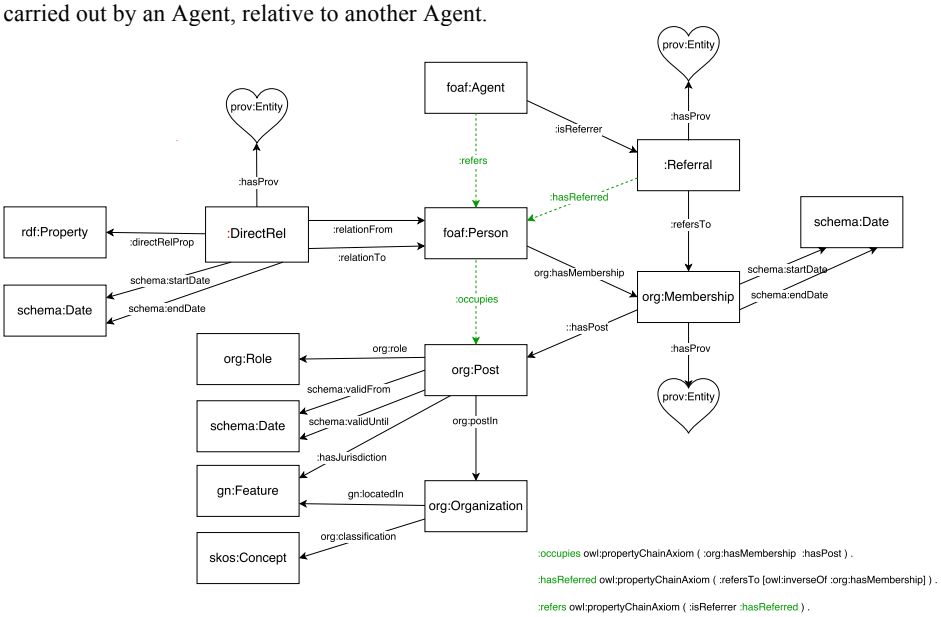

Figure 6 – The complete POLARE ontology including provenance information.

Figure 6 shows the complete POLARE ontology, where the heart symbol represents the PROVHeart pattern. One can see that claims involving any of the main relations (Membership, Referral and DirectRel) have provenance information attached to them.

As an example of PROVHeart use, consider the case where someone wants to know all Persons, and respective Posts, that were referred by a Person called "Jader Barbalho". Furthermore, these Referrals should have provenances stating that the sources used (evidence) were published by a Person called "Andréia Sadi" or are from a Magazine called "Época", and, also, should be published to the database by a Person called "Daniel" or "Laufer".

We implemented claims based on the nanopublication model [12][13], where a claim, represented by a set of statements, is recorded as a group of three named graphs:

- assertion graph - containing the set of statements that compose the claim, using POLARE ontology;
- publication info - containing statements that give provenance information about the claim publication activity, using the provHeart pattern;
- provenance - containing statements that give provenance information about the claim itself, using the provHeart pattern.

In our implementation, we used a named graph, called claims, to store the nanopublications related to the set of all claims published in the database.

The example above can be queried in the database using the following query:

[13] http://www.nanopub.org

```
SELECT DISTINCT ?referrerName ?referredName ?roleName ?startdate ?citation
WHERE {
 GRAPH graph:claims {
  ?pub np:hasAssertion ?claim ;np:hasProvenance ?prov ;np:hasPublicationInfo ?pubinfo . }
 GRAPH ?claim {
  ?referral a polare:Referral ;polare:hasReferred ?referred ;schema:startDate ?startdate .
  ?referrer polare:isReferrer ?referral ;foaf:name ?referrerName .
  ?referred polare:occupies ?post ; foaf:name ?refferedName .  ?post org:role ?role .
  ?role skos:prefLabel ?roleName .
  FILTER (?referrer = ex:JaderBarbalho) . }
 GRAPH ?prov {
  ?claim prov:hadPrimarySource ?source ; prov:wasAttributedTo ?agent ;
         prov:wasGeneratedBy ?activity .
  ?source dcterms:bibliographicCitation ?citation .
  ?delegation prov:hadActivity ?activity ;prov:agent ?org .
  FILTER (?agent = ex:AndreiaSadi || ?org = ex:EpocaMagazine) . }
 GRAPH ?pubinfo {
  ?pub prov:wasAttributedTo ?agpub .
  FILTER (?agpub = ex:Daniel || ?agpub = ex:Laufer) . }
```

In terms of the Trust Framework, the claims discussed here allow describing Data/Information items and Provenance information about them. Further discussion about trust and trust policies is outside the scope of this paper.

4 POLARE in action

In this section, we present an example of how the database can be used to better understand a news story, discussing the options made available to the reader to establish her/his trust on the information contained in the story.

Consider the blog post[14] made by reporter Andréia Sadi, who works for the GloboNews news channel, about an investigation carried out by the Federal Police in the office of Congresswoman Simone Morgado, wife of Senator Jader Barbalho. This investigation is related to embezzlement accusations against an aide working in the Congressswoman's office. Figure 7 contains a simplified diagram with the main facts regarding the Persons, Posts, and Organizations cited in the blog post. In this diagram, we use a simplified view of provenance, showing only a few provenance facts using the "source" link.

The information in the SNLP database contains the following information about the persons involved in the blog post.

Jader Barbalho is a Brazilian lawyer, politician and businessman. In his long political career, he was a Federal Deputy of Pará, Governor of Pará, Minister of Social Security of Brazil and Minister for Agrarian Development of Brasil. He occupies the post of Senator of Pará in the current legislature. He was married to Simone Morgado, until the end of 2016, according to his son, Helder Barbalho.

Helder Barbalho is a Brazilian politician and administrator. He was, during 8 years, the Mayor of Ananindeua, a city located in the state of Pará. He was the former Minister of Fisheries and Aquaculture of Brazil and is currently the Minister of National Integration.

Simone Morgado is a Brazilian economist and politician. She was Mayor of Bragança, a city located in the state of Pará and is currently a Congresswoman representing the state of Pará.

Soane Castro currently occupies the post of aide in the office of Simone Morgado, and held the post of Superintendent in the Ministry of Fisheries and Aquaculture of Brazil, during the period when Helder Barbalho was its Minister.

[14] http://g1.globo.com/politica/blog/andreia-sadi/post/pf-realiza-operacao-na-camara-dos-deputados-mas-alvo-nao-e-localizado.html. (In Portuguese)

The shaded area indicates facts (actually, claims) that are not directly cited in the story, but are part of the database. In particular, it explicates that there are several connections linking Senator Jader Barbalho to the aide who is the target of the investigation – she was appointed for the position related to the embezzlement accusations by the Senator's son, Helder Barbalho, who was the Minister in charge of that agency at that time. This post (position) is in the home state of all agents involved, Pará.

Let us focus now on the kinship information about Simone Morgado and Jader Barbalho. The blog post itself states that they are married, but also adds that Helder Barbalho claims they have not been married since November 2016. The fact that the reporter did not directly refer to Ms. Morgado as "Senator Jader Barbalho's ex-wife" seems to indicate she chose not to fully trust the information that they are separated, and instead refers the reader to a secondary source citing the Senator's son claim about this fact. However, she did not give any provenance to Helder Barbalho's claim – i.e., how she became aware of the claim. Figure 8 shows the POLARE instance for the kinship relation claimed in the blog post, including reference to PROVHeart instances, Prov_2 and Prov_3.

Assuming that agent Laufer was responsible for entering the facts mentioned in the blog post into the database in wiki style, the corresponding provenance information Prov_2 using PROVHeart is shown in Figure 9.

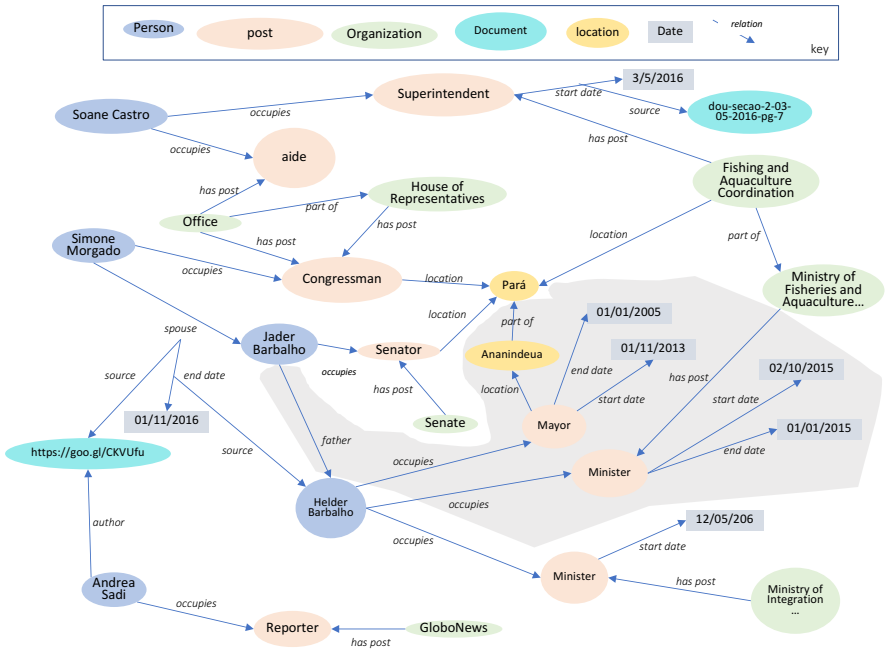

Figure 7 – An instance of POLARE with facts underlying a blog post

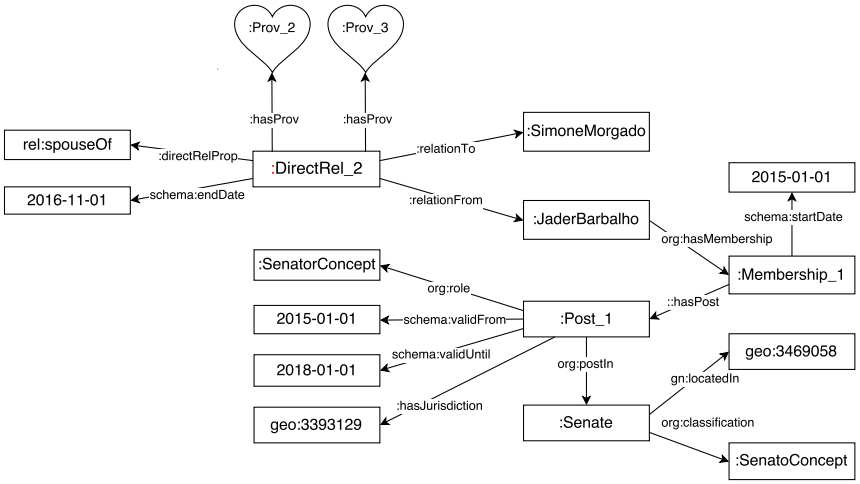

Figure 8 – Kinship relation modeled using POLARE

In this provenance information, Laufer has chosen to interpret the blog post as stating that the spouse relationship has an end date, in spite of the slight "conflict" in the text. It would be possible to create a separate claim for the end date should this be considered relevant and important.

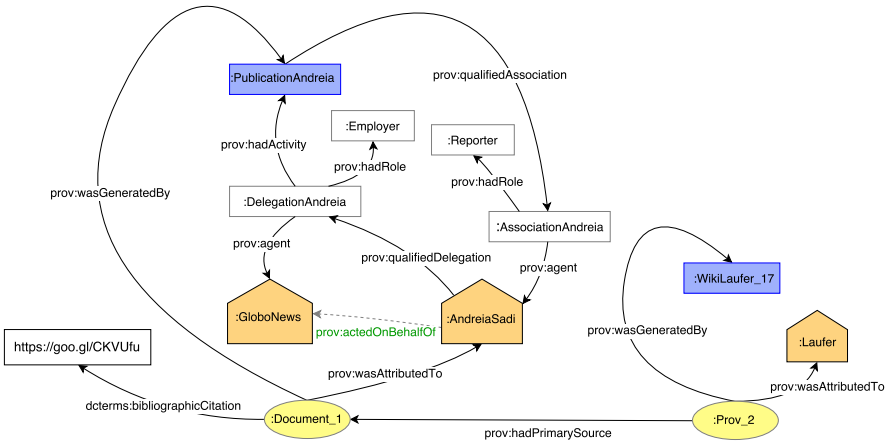

Figure 9 - First provenance information about blog post (Prov2).

Suppose now that another agent, Daniel, decided to check whether this kinship information is accurate. Short of trying to find a court document stating that they are officially divorced, Daniel searches for additional sources for this claim, and finds a news story stating the same fact, and enters it into the database. This becomes a second provenance information for the claim in the database, which is represented in Figure 8 as the Prov_3 PROVHeart instance.

As mentioned, the actual trust process is outside the scope of this paper, but we nevertheless outline how the provenance information can be used in a trust process. In the trust process, the consumer of this data must apply her/his own policies in order to accept the truth of the statement.

In the example, one policy could be to simply accept it (meaning, to use it in her/his computations or decisions) based on her/his trust on the recording agent, e.g., Laufer. A second policy, a bit more cautious, would be to accept it based on the agent responsible for the primary source, Andréia Sadi. At this point, it is not possible to continue the chain of provenance, because there is no additional provenance information given in the blog post. A different kind of policy can be used, however. Analogously to journalistic principles of obtaining multiple independent sources about the truth of some claim, the consumer could decide to trust this claim based on the fact that there are two different provenances given, although, strictly speaking, s/he should also check whether Agents Andréia Sadi and Murilo Ramos are independent.

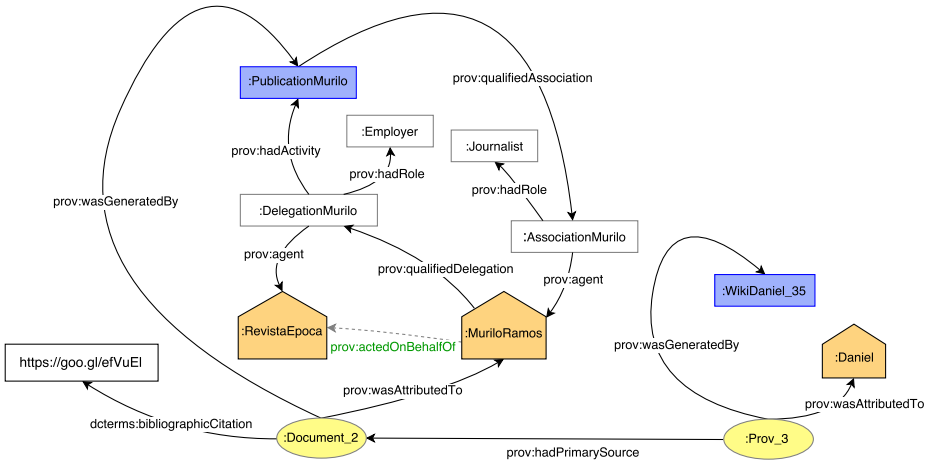

Figure 10 - Second provenance information about blog post (Prov3).

5 Related Work

Our work is related to several others along different aspects. With respect to the project itself, and its goals of making connections explicit, we can cite the LittleSis.org[15] website, and the Poderopedia Project[16]. Both have built databases about Political Agents, the former specifically for the US, the latter allowing instantiations for different countries. Compared to SNLP (and POLARE) LittleSis has a rich model for organizations, and similar types of relations between People and Organizations. Poderopedia, in addition, has a finer grained model for some of the relations. None of them include the Referral relation.

Whereas both allow representing provenance, they do so by linking documents as sources, and therefore it is not possible to use chained provenance data in the way enabled by POLARE. whose central characteristic is the inclusion of patterns of provenance information, as an extension of data itself. Since both projects allow inclusion by users, they have a much stricter and structured editorial process to ensure the quality of the information and uniform use of criteria, a task that is more easily automated by the use of the PROVHeart pattern. As discussed, this finer grained provenance information also provides better support for the trust process.

[15] https://littlesis.org/
[16] http://www.poderopedia.org/poderopedia/pages/index/2

The Popolo Project[17] is an initiative to define data interchange formats and data models for governments, in the context of Open Government. They define a set of classes that cover, basically, the specification of Persons and Organizations, using Posts and Memberships to relate Persons to Organizations. Popolo also defines a set of classes related to voting processes. The main principle of Popolo is reuse. The classes are defined as a set of properties of well-known vocabularies, including ORG, FOAF, geonames, schema.org, etc. POLARE defines the linking between Persons to Organizations in a similar way. POLARE introduces a set of relations that are not contemplated by Popolo, as kinship relations and referrals. Popolo does not include provenance information.

6 Conclusions and Future Work

We have presented a framework for describing the Trust Process, and used it to contextualize the use of Semantic Data in the SLNP project, which aims at building an open Linked Data database of Political Agents, using the POLARE ontology. Special attention is given to provenance data in order to support a trust process.

There are several directions in which we are continuing this research. We have already made two additional ontologies to represent the Voting Process in the Legislative branch and the Electoral Campaign Process. The former allows recording the various bills proposed in Congress, and how Congressmen voted on them throughout the legislative process. The latter allows recording every candidate's income declaration, the contributions to their campaign as well as payments made by campaign offices to suppliers and to other campaigns. It is our goal to also include information about organizations in a manner similar to LittleSis as specializations to the Org ontology, and by adding more direct relations between Organizations.

Another important type of relation between Agents are transactions; we plan to enrich the POLARE ontology to include them. In our view, transactions should include business transactions, but also legal actions, educational actions, etc.…

Whereas PROVHeart is able to capture the more common provenance pattern found in the Political Agents domain, it is possible that more detailed or finer grained patterns are necessary, something we are also investigating, and plan to evaluate it with journalists involved in the project.

Even with provenance information, there is still an issue of authentication of documents – they could be false after all. This falls under the realm of digital signatures, and this should be integrated into the framework. In addition, the provenance information itself could be falsified. There are emerging technologies such as blockchains that may be used to address this, providing a verification that the provenance chain has not been tampered with, so long as each provenance assertion is digitally signed as well.

With respect to implementing provenance information, the approach described here uses the traditional reification via class approach. We have also been investigating the idea of reifying the properties using the singleton property approach [21], which allows modeling properties (actually property instances) as individuals and as Properties. Figure 12 shows a snippet of the POLARE schema using this approach, and the corresponding OWL fragment of an instance. We want to investigate if the differences in these approaches could reveal a better way of constructing queries to inspect data in the database.

Another implementation aspect we will investigate is the addition of access control to portions of the database. We envisage that certain communities (e.g., reporters) may want to enter information to the database but restrict, at least for a period of time, general access to it.

[17] http://www.popoloproject.com/

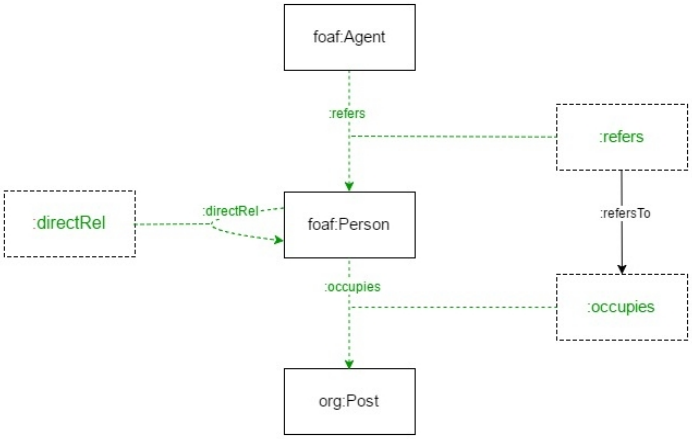

Figure 11 - Representing properties using the Singleton approach.

```
:JaderBarbalho rdf:type owl:NamedIndividual ,foaf:Person ;
            foaf:name "Jader Barbalho"^^xsd:string ;
            :occupies_1 :Post_1 ;
:occupies_1 rdf:type owl:NamedIndividual, owl:ObjectProperty,
                  org:Membership ;
                 schema:startDate "2015-01-01"^^xsd:date ;
                 :singletonPropertyOf :occupies.
:occupies rdf:type owl:NamedIndividual, owl:ObjectProperty.
```

Figure 12 – Representing relations as singletons

7 References

1. Shorenstein Center on Media, Politics and Public Policy, Harvard Kennedy School, "Fake News and the Spread of Misinformation", , https://journalistsresource.org/studies/society/internet/fake-news-conspiracy-theories-journalism-research, accessed May 13, 2017.
2. Albright, J.; "Welcome to Fake News", https://medium.com/@d1gi/election2016-fakenews-compilation-455870d04bb, accessed May 13, 2017
3. Allcot, H.; Gentskow, M.; Social Media and Fake News in the 2016 Election. Journal of Economic Perspectives 31(2),Spring 2017, pp 211–236
4. Almendra, V. D. S., & Schwabe, D. (2006). Trust policies for semantic web repositories. In *Proceedings of 2nd International Semantic Web Policy Workshop (SWPW'06), at the 5th International Semantic Web Conference (ISWC 2006)* (pp. 17-31).
5. Artz, D; Gil, Y.; A survey of trust in computer science and the Semantic Web, Web Semantics: Science, Services and Agents on the World Wide Web, Volume 5, Issue 2, 2007, Pages 58-71, ISSN 1570-8268, http://dx.doi.org/10.1016/j.websem.2007.03.002.(http://www.sciencedirect.com/science/article/pii/S1570826807000133)
6. Baker, T.; "Fake News: What Is It, And How Can We Tackle It? https://digitalsocial.eu/blog/58/fake-news-what-is-it-and-how-can-we-tackle-it, accessed May 13, 2017.

7. Bizer, C.; and Oldakowski, O;.. Using context- and content-based trust policies on the semantic web. In *Proceedings of the 13th international World Wide Web conference on Alternate track papers & posters* (WWW Alt. '04). ACM, New York, NY, USA, 228-229. 2004 DOI: https://doi.org/10.1145/1013367.1013409
8. Ciampaglia, G. L.; Shiralkar, P.; Rocha, L. M.; Bollen, J.; Menczer, F.; Flammini, A.; "Computational Fact Checking from Knowledge Networks", PLoS ONE, 10(6): e0128193. 2015. http://doi.org/10.1371/journal.pone.0128193
9. Conroy, N. J., Rubin, V. L. and Chen, Y. (2015), Automatic deception detection: Methods for finding fake news. Proc. Assoc. Info. Sci. Tech., 52: 1–4. doi:10.1002/pra2.2015.145052010082
10. Dutton, W. H. (2009). The Fifth Estate Emerging through the Network of Networks. Prometheus, 27(1), 1–15. http://doi.org/10.1080/08109020802657453
11. Gil, Y.; Artz, D.; Towards content trust of web resources. *Web Semant.* 5, 4 (December 2007), 227-239. DOI:http://dx.doi.org/10.1016/j.websem.2007.09.005
12. Groth, Paul, Andrew Gibson, and Jan Velterop. "The anatomy of a nanopublication." *Information Services & Use* 30.1-2 (2010): 51-56.
13. Jarvis, J.; Borthwick. J,."A Call for Cooperation Against Fake News", Medium, https://medium.com/whither-news/a-call-for-cooperation-against-fake-news-d7d94bb6e0d4, accessed May 13, 2017.
14. Angella J. ;Kim, K; K.P. Johnson, Power of consumers using social media: Examining the influences of brand-related user-generated content on Facebook, Computers in Human Behavior, Volume 58, 2016, Pages 98-108, ISSN 0747-5632, http://dx.doi.org/10.1016/j.chb.2015.12.047. (http://www.sciencedirect.com/science/article/pii/S0747563215303186)
15. Lehrman, S. (ed), "Trust Project Summit Report", Markkula Center for Applied Ethics, Santa Clara University, 2016. https://www.scu.edu/ethics/focus-areas/journalism-ethics/programs/the-trust-project/trust-project-summit-report/. See also Trust Project Indicators, https://trello.com/b/YbHYmodO/trust-project-indicators.
16. Mayer, R. C., Davis, J. H., & Schoorman, F. D.. An integrative model of organizational trust. Academy of management review, 20(3), 709-734, 1995.
17. McKnight, D. H., & Chervany, N. L. (2001). Trust and distrust definitions: One bite at a time. In Trust in Cyber-societies (pp. 27-54). Springer Berlin Heidelberg.
18. Mele, N., Lazer, D., Baum, M., Grinberg, N., Friedland, L., Joseph, K., & Mattsson, C. (2017). Combating Fake News: An Agenda for Research and Action., http://www.northeastern.edu/nulab/wp-content/uploads/2017/04/Combating-Fake-News-Agenda-for-Research.pdf
19. Momeni, E., Cardie C., and Diakopoulos, N.. A Survey on Assessment and Ranking Methodologies for User-Generated Content on the Web. *ACM Comput. Surv.* 48, 3, Article 41 (December 2015), 49 pages. DOI: http://dx.doi.org/10.1145/2811282
20. Moreau, L., Batlajery, B., Huynh, T. D., Michaelides, D., & Packer, H. (2017). A Templating System to Generate Provenance. Accepted for publication, *IEEE Transactions on Software Engineering*. http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7909036
21. Nguyen, V., Bodenreider, O., & Sheth, A. (2014, April). Don't like RDF reification?: making statements about statements using singleton property. In *Proceedings of the 23rd international conference on World wide web* (pp. 759-770). ACM.
22. Pinyol, I.; Sabater-Mir, J.; Computational trust and reputation models for open multi-agent systems: a review. Artificial Intelligence Review 40:1–25 DOI 10.1007/s10462-011-9277-z (2013)

23. Ruchansky, N., Seo, S., & Liu, Y. (2017). CSI: A Hybrid Deep Model for Fake News. arXiv preprint arXiv:1703.06959.
24. Sherchan, W., Nepal, S., and Paris, C. 2013. A Survey of trust in social networks. ACM Comput. Surv. 45, 4, Article 47 (August 2013), 33 pages. DOI: http://dx.doi.org/10.1145/2501654.2501661
25. Shao, C.; Ciampaglia, G. L.; Flammini, A.; and Menczer, F.; Hoaxy: A Platform for Tracking Online Misinformation.. In Proc. Third Workshop on Social News on the Web (WWW SNOW), 2016. Preprint arXiv:1603.01511 http://doi.org/10.1145/2872518.2890098
26. Starbird, K.; "Examining the Alternative Media Ecosystem through the Production of Alternative Narratives of Mass Shooting Events on Twitter", The 11th International AAAI Conference on Web and Social Media (ICWSM-17), Montreal, CA, 2017 pre-print http://faculty.washington.edu/kstarbi/Alt_Narratives_ICWSM17-CameraReady.pdf
27. Starbird, K.; "Information Wars: A Window into the Alternative Media Ecosystem", https://medium.com/hci-design-at-uw/information-wars-a-window-into-the-alternative-media-ecosystem-a1347f32fd8f, accessed on May 12, 2017
28. Wang, W. Y. (2017). " Liar, Liar Pants on Fire": A New Benchmark Dataset for Fake News Detection. arXiv preprint arXiv:1705.00648.
29. Wardle, C., "Fake News. It's Complicated", FirstDraft News, https://firstdraftnews.com/fake-news-complicated/, accessed on May 12, 2017.