

The eLOD Ontology: Modeling Economic Open Data

Michalis Vafopoulos, Gerasimos Razis, Ioannis Anagnostopoulos, Georgios Vafeiadis, Dimitrios Negkas, Eleftherios Galanos, Aggelos Tzani, Ilias Skaros, Konstantinos Glykos

► **To cite this version:**

Michalis Vafopoulos, Gerasimos Razis, Ioannis Anagnostopoulos, Georgios Vafeiadis, Dimitrios Negkas, et al.. The eLOD Ontology: Modeling Economic Open Data. 12th IFIP International Conference on Artificial Intelligence Applications and Innovations (AIAI), Sep 2016, Thessaloniki, Greece. pp.119-130, 10.1007/978-3-319-44944-9_11 . hal-01557596

HAL Id: hal-01557596

<https://hal.inria.fr/hal-01557596>

Submitted on 6 Jul 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



The eLOD Ontology: modeling economic open data

M. Vafopoulos¹, G. Razis², I. Anagnostopoulos², G. Vafeiadis¹, D. Negkas³, E. Galanos³, A. Tzani³, I. Skaros³, K. Glykos³

¹Software and Knowledge Engineering Laboratory, IIT, NCSR-“Demokritos”
Athens, Greece
E-mail: vaf@aegean.gr, vafeiadis.giorgos@gmail.com

²School of Sciences, Dpt of Computer Science and Biomedical Informatics
University of Thessaly, Lamia, Greece
E-mail: {razis, janag}@dib.uth.gr

³University of Piraeus Research Centre, Piraeus, Greece
E-mail: linkedeconomy@gmail.com

Abstract. For decades, valuable economic data was out of reach for most of the people. Gradually, public budgets and tenders are becoming openly available and global initiatives promote financial transparency and data innovation. But, yet the poor quality of open data undermines their potential to answer interesting questions (e.g. efficiency of public funds and market processes). The eLOD ontology has been initiated as a top-level conceptualization that interlinks the publicly available economic open data by modelling the flows incorporated in public procurement together with the market process to address complex policy issues. This paper presents the basic aspects of eLOD ontology in interlinking and querying diverse open data ranging from budget execution to prices. Already, eLOD ontology is used by two EU projects in order to develop new systems, to enable information exchange between systems, to integrate data from heterogeneous sources and to publish open data related to economic activities.

Keywords: Linked Data, Semantic Web, Knowledge Representation, Economy.

1 Introduction – Related works

In this paper, we introduce the eLOD Ontological schema as the “heart” and basic data modelling and handling infrastructure of LinkedEconomy¹ portal. LinkedEconomy aims at providing a universal access to Greek and international economy data, as well as at promoting the benefits of linking heterogeneous sources under the concept of open and reusable data in this critical domain [1]. In addition, it targets to offer a unique platform of rich linked economy data for enhancing the citizens’ awareness in

¹ <http://linkedeconomy.org/en/>

respect to economic issues in Greece and worldwide, as well as to provide curated and semantified data to the Linked Open Data research community. It is the ancestor of the Public Spending (PSNET) initiative², which was the first attempt to harvest, align, interlink, analyze and distribute as Linked Open Data massive amounts of public spending data from Greece and from six other (local and national) governments [2]. Related worldwide initiatives are driven from similar Open Government Data projects. Below, we classify them in four major fields according to their applications in the economy domain.

In the field of open budgets, the International Budget Partnership advocates for public access to accountable budget systems. One of its projects, the Open Budget Survey Tracker³ allows citizens to monitor whether central governments are releasing the requisite information on how the government is managing public finances. Starting from 2015, EU funds the Open Budgets project⁴ to provide a scalable platform for public administrations to publish open budget data that is easy-to-use, flexible, and attractive for all. In addition, the Open Spending Initiative⁵ offers an easy system to upload, explore and share public finance data (e.g. budgets or expenditure databases). In 2013, more than 16 million transactions from 363 different datasets around the world have been uploaded. A similar project from the Sunlight Foundation⁶ analyzes the spending data uploaded in the official portal of US government. OpenTED⁷ is an initiative that provides data dumps of tenders from the joint European procurement system more easily accessible to journalists and researchers.

In the field of data standards, the Fiscal Data Package⁸ is developed as a simple, open technical specification for government budget and spending data. Public Contracts Ontology has been introduced by the LOD2 project [3] to provide an ontological basis for representing key concepts in tenders and expressing structured data about public contracts. The Core vocabularies have been developed by European Commission's Interoperability Solutions for European Public Administrations (ISA) programme as simplified, generic, re-usable and extensible data standards that model the characteristics of basic entities in a context-neutral and syntax-neutral fashion. Most of them have been incorporated in our modelling, as it will be described below. Open Contracting Partnership⁹ is a consultation process to create a set of global principles that can serve as a guide to advance open contracting around the world. In this context, the Open Contracting Data Standard (OCDS) sets out key documents and data that should be published at each stage of a contracting process.

In the field of company data, Open Corporates¹⁰ aggregates company information from different countries and jurisdictions. The specific team is working on creating

² <http://publicspending.net>

³ <http://obstracker.org>

⁴ <http://openbudgets.eu>

⁵ <http://openspending.org>

⁶ <http://clearspending.org>

⁷ <http://ted.openspending.org/>

⁸ <http://fiscal.dataprotocols.org>

⁹ <http://www.open-contracting.org/>

¹⁰ <http://opencorporates.com>

Linked Data representations out of their databases, by mapping company metadata to certified ontologies such as the Core Business Vocabulary and linking them to other data hubs, such as DBpedia.org and Geonames. Financial statements and accounting reports that are published by companies contain important data.

In the area of product data, Open Product Data - POD¹¹ hosted by OKF as a community project, stands as a public database of product data connected to barcodes, in order to empower consumers with useful and machine-readable product information (e.g. prices).

Other related research approaches have a two-fold orientation. In the first fold, we have efforts that analyse the benefits of using open data in economy, trying in parallel to provide a unified framework and modelling strategy [4], [5], [6], while, in the second fold, we have the contributions that integrate the benefits of openness, thus creating a flourishing linked ecosystem of economy data [7], [8], [9], [10].

2 eLOD ontology: modelling economic data under semantics

In this section, we describe some well-known and established structured vocabularies, which are used in the economy domain, as well as the model that incorporates interconnections between public finances and market processes in Greece and consists of the main classes and properties of our ontology.

DATA SOURCES USED		
Name	Provision through API	URI
Transparency	yes	diavgeia.gov.gr
Central Electronic Registry of Public Procurement	no	www.eprocurement.gov.gr
National Strategic Reference Framework	no	www.espa.gr/en
Central Market of Thessaloniki	no	www.kath.gr
Prices Observatory	yes	www.e-prices.gr
Fuel Prices Observatory	yes	www.fuelprices.gr
Municipality of Athens	no	www.cityofathens.gr/khe/proypologismos
Municipality of Thessaloniki	no	http://www.thessaloniki.gr

Table 1. Related information for Data Sources used

¹¹ <http://product.okfn.org>

2.1 Description of sources and vocabularies used

Table 1 depicts the data sources used as well as some information regarding them (API support, URI). During the ontology design phase, our basic aim was to reuse other well-known and established vocabularies and ontologies that cover our needs. When existing models were inadequate we introduced and defined new concepts. The well-established ontologies and vocabularies we used are FOAF¹², GoodRelations¹³, Public Contracts¹⁴, Organization Ontology¹⁵, Registered Organization Vocabulary¹⁶, Dublin Core¹⁷, SKOS¹⁸ and vCard¹⁹, while we also use some properties from the DBpedia Ontology²⁰.

FOAF, which is an acronym of “Friend of a Friend”, is a vocabulary for describing people, their activities and relationships with other people and objects. This concept can be generalized in order to describe all types of entities, named “agents”, who are responsible for specific actions. FOAF aims to describe the world by using simple ideas inspired by the Web. In our economic context, this vocabulary is used in order to define and describe the agents who are responsible for specific actions. Two categories of agents are used, namely “Persons” and “Organizations”.

GoodRelations is an ontology, which aims at defining a data structure for products related to electronic commerce, the prices, the stores and the data of the companies. Its use allows the expression of the commercial and operational details of scenarios for e-commerce. The main entities in this domain are the involved agents, in terms of persons or companies, the objects of the commercial activities, the items for sale, lease or repair, as well as the locations where such offers are available. In our economic context, this ontology is used in order to identify and describe the “Business Entities” which are involved in a commercial activity (their legal names and their Vat Ids), the type of their services, and the financial details of the contract or of the payment (i.e. the wholesale or retail price, whether tax is included in the price and the expressed currency of the price). Moreover, it is also used to express a point or area of interest from which a particular product or service is available.

Public Contracts is an ontology that aims at describing the contracts in the public sector. It is based on the “GoodRelations” ontology for the modeling of business entities and price specifications.

In our economic context, this ontology is used in order to define and describe the following:

¹² <http://www.foaf-project.org>

¹³ <http://purl.org/goodrelations>

¹⁴ <https://code.google.com/p/public-contracts-ontology>

¹⁵ <http://www.w3.org/TR/vocab-org>

¹⁶ <http://www.w3.org/TR/vocab-regorg>

¹⁷ <http://dublincore.org>

¹⁸ <http://www.w3.org/2004/02/skos>

¹⁹ <http://www.w3.org/TR/vcard-rdf>

²⁰ <http://wiki.dbpedia.org/services-resources/ontology>

- the public contracts during all stages of their existence (tender, contract, payment),
- the procedures specifying how the details of a contract are published and how a supplier is chosen,
- the basic focus of the contract (i.e. works, supplies or services),
- the price of the contract, depending on its phase (before or after the offer),
- the award criteria that define the conditions under which the best offer will be selected and awarded along with their weights, and
- the main and supplementary products or services purchased by the contract (determined by the CPV codes).

The *Organization* ontology aims at describing the organizational structures in order to support the disclosure interconnected data of organizational information in various sectors. It is designed as to allow extensions in specific areas for classifying the organizations and roles as well as for extensions in order to support relevant information, such as organizational activities. Its design allows the publication of information on organizations and their organizational structures, including governmental ones. In our economic context, this ontology is used in order to define and describe the organizations and their organizational units. Their structure is also represented using the properties of this ontology. Additional properties are also provided in order to illustrate the members and their structures within an organization, as well as the roles, positions and relationships between people and organizations.

The *Registered Organization* vocabulary aims at describing the entities, which have obtained the status of a legal entity through a formal registration process, usually at the national or regional registry. It includes a minimum number of classes and properties that are designed to depict the typical information recorded by the business' registries, thus facilitating the exchange of information between them, despite having considerable variation between the recorded and the published data. In our economic context, this vocabulary is used in order to identify and describe the business entities and their properties, including their type, status and activity.

Dublin Core and *Dublin Core Terms* is a small vocabulary for the description of general metadata of the Web and of natural resources. In our economic context, this vocabulary is used in order to define the entities that are responsible for publishing a contract, the subject of the contract, and the date of its formal issue.

SKOS, which is the acronym of "Simple Knowledge Organization System", is a vocabulary, which is designed in order to represent thesaurus, classification schemes, classifications, lists or any other type of structured controlled vocabulary. Its main aim is to allow the easy publication and use of these vocabularies as interconnected data. In our economic context, this dictionary is used in order to define numerous controlled vocabularies and to represent code lists (e.g. Currencies).

The specification of *vCard* is generally used for describing people and organizations. Usually, *vCard* objects are encoded based on their own syntax or in their XML format. The *vCard* ontology can be used for the semantic representation of any *vCard* data and is also focused on describing people and organizations, including location information and groups of such entities, as "FOAF" and "Organization" do. In our

economic context this ontology is used in order to represent geographical information regarding agents and products, such as delivery addresses.

DBpedia is an ontology, which covers different domains and is created manually based on the most frequently used information from infoboxes of Wikipedia. For the purposes of our case, properties of this ontology are used in order to define different type of information about businesses (e.g. subsidiaries of a company).

2.2 The eLOD ontological schema

The ontology of each data source along with its description and example queries can be found on Github²¹. In this repository, we describe in an analytical way the ontological schema created for the modeling of each data source described in Section 2. Apart from the ontology description, SPARQL queries are provided along with a sample of their responses. The model of our ontological schema is depicted in Figure 1. As it shows, government forms and publishes budgets, parts of which include projects and works that are assigned through calls for tenders. After contracts have being signed and projects are fulfilled, funds are transferred. Spending data are often used to assess the completion of public budgets. However, another type of added-value fiscal information that has been started to provided publicly is subsidies and aid data. Subsidies include government payments to firms and households based on a development plan (e.g. in Greece we have the Greek National Strategic Reference Framework (anaptyxi.gov.gr/), while in EU the farmsubsidy.openspending.org). Procurements, subsidies and aid awarding processes are involving the exchange of information among authorities (e.g. tax offices, business registries and various public agencies) and the official publication of relevant information (e.g. call for tenders, payments).




Fig. 1. Ontological model that incorporates interconnections between public finances and market processes

²¹ <https://github.com/LinkedEcon/LinkedEconomyOntology-ELOD>

In addition to the above, we have also modeled the Market domain. Information about the market process (e.g. price, value of sales and quality) is an important source of business value and is mostly closed into corporate environments. Market process information is partially shared with government authorities (e.g. tax office), suppliers and consumers. Although, there is a small but crucial part of business information that should be shared publicly as open data to ensure that quality and competition are best served. This set of publicly available information should include the vector of prices (at least for basic goods and services), aggregated quantities sold in wholesale and retail markets and all the relevant input to assess quality of provisioned products and services. In Greece, a representative sampling of retail prices for thousands of basic consumer goods and fuels can be retrieved on a regular basis. In some cases, (e.g. Central Market of Thessaloniki) wholesale prices and quantities are provided for fruit, vegetables and meat. In order to make clearer to the reader the reuse and incorporation of the aforementioned ontologies and vocabularies to the eLOD Ontology, two informative tables are provided. Table 2 presents the percentage of the classes and properties belonging to each vocabulary.

Vocabulary	Distribution of Classes in eLOD Ontology (%)	Distribution of Properties in eLOD ontology
{own semantics}	50%	67.7%
Good Relations	15.9%	8.7%
Public Contracts	9.1%	7.7%
Organizations	9.1%	3.1%
SKOS	4.5%	3.1%
Registered Organization	4.5%	2.6%
FOAF	4.5%	1.0%
vCard	2.3%	2.6%
Dublin Core	-	3.1%
DBpedia	-	0.5%

Table 2. Classes / Properties percentage use of established vocabularies in eLOD Ontology

2.3 Approach and reuse

The eLOD ontology is based more on a theory-driven than in data- or statistics-driven approach. It has been designed to better balance the trade-off of being as generic as to be scalable to future open data categories and as specific as to be compatible with existing initiatives. It could be an opportunity for many of the diverse communities, which are working on transparency, global standards and economic data, to join forces in addressing useful and not-yet-answered questions. For example, we can't yet answer, if public spending is expensive and comparable across countries? or

can we compare financial ratios in public budgets? Can we comparatively analyze wholesale and retail prices? Due to its theory-driven approach eLOD is used by two major European projects related to open economic data. Your Data Stories (YDS, yourdatastories.eu) is an EU funded project (Grant Agreement No.645886) that aims to convert publically available economic open data into re-usable and interoperable building blocks that can be used to construct applications. YDS will allow any actor to design and implement personalised public services. Based on the eLOD model feeds from trusted sources will be interconnected with the new and re-purposed data feeds provided by users of the social web in order to form a meaningful, searchable, customizable, re-usable and open hyper-market of data, feeds, and services. The second project is Big Data Europe (<http://www.big-data-europe.eu/>) which is funded by EU (grant agreement n. 644564) for enabling European companies to build innovative multilingual products and services based on semantically interoperable, large-scale, multi-lingual data assets and knowledge, available under a variety of licenses and business models. LOE model has been selected to support the implementation of one of the seven societal challenges that Big Data Europe focuses. In particular, the social sciences challenge refers to statistical and research data linking and integration and will focus on citizens budgeting and control.

Diavgeia	Number of Payments	11,288
	Total Amount (in Euros)	44,550,925.34
e-Procurement	Number of Payments	289
	Total Amount (in Euros)	30,570,036.85
NSRF	Number of Subsidies	5
	Total Amount (in Euros)	208,311

Table 3. Query 1 results [Hellenic Telecommunications Organization S.A.]

Expense Approval Counter	7,738
Expense Approval Amount	35,746,746.58
Founding Year	1949
Net Income	676,400,000
Number of Employees	29,330

Table 4. Federated Query 2 results [Hellenic Telecommunications Organization S.A.]

3 Asking the data: a case study

The aim of LinkedEconomy.org is not only to transform and semantically enrich the input data into RDF graphs, but also to apply a unified ontological model in order to treat and query distinct datasets as one.

As already mentioned, a public endpoint²² allows the search of the collected semantic data and their combination and enrichment with other sources. SPARQL Query 1 displays the advantages of having a unified ontological model, as it combines such data from three different sources, namely Diavgeia, e-Procurement and NSRF. The query returns the number of decisions and subsidies referring to payments and their total amount in Euros for the case of the Hellenic Telecommunications Organization S.A (member of the Deutsche Telekom Group). For these decisions this organization appears as “seller”, meaning that it receives payments for offering a service or from selling a product.

One of the main characteristics of the semantic data is that other publicly available datasets from the LOD cloud can at the same time be queried and combined with relevant information. Such an example is the federated SPARQL Query 2, which combines the economical data of Diavgeia, as stored in our graph, along with DBpedia information. Table 2 depicts the results we receive when we “ask” our data combined with DBpedia for the case of the Hellenic Telecommunications Organization S.A. Apart from the information from Diavgeia, the user receives information in respect to the organization’s founding year, net income and number of employees.

4 Discussion - Future work

This paper presents i) the economic open data sources used, ii) the ontological model that orchestrate linked data flows, iii) the common controlled vocabularies that are incorporated in eLOD ontology in order to place it the global economic Linked Data cloud. All consist of the “data engine” of LinkedEconomy, a project that lead to publishing data of high added value not only for exploration purposes, but also for exploiting the benefits of transparency and openness to the citizens, the research community, and the government itself. Through this research initiative, our team tackled the challenge of building a common terminology for the basic financial and economy activities, which will -in turn- facilitate the research over new linked data and sources. All described components form a system capable of linking economy-related data at large scale, creating in parallel a framework for collect, validate, clean and publish linked data streams. Our efforts produced a CKAN repository, which publishes datasets from sources that are being updated regularly and contain valuable information in respect to social and economic research²³. Citizens and economy stakeholders (government, local authorities, etc.) can exploit 27 datasets from 14 classified data sources in machine-readable format (xlsx, csv, rdf). Examples include economy data,

²² <http://elod7.linkedeconomy.org:8890/sparql>

²³ <http://ckan.linkedeconomy.org>

such as public procurements, budgets, prices, expenditures, data from the insurance domain and employment, as well as financial and macro-economic data. As we envisage the use of Web 3.0 technologies and we acknowledge the benefits of openness, we offer a publicly available SPARQL endpoint that consists of more than 210M triplets in total, and we share all ontological schemas and related specifications in Github.

Finally, as publicly available open data are growing rapidly worldwide, we currently work on modeling many foreign economic datasets according the already semantic knowledge we developed. We expect to support really soon the provision of more than 15 different economic datasets from Europe, UK, Australia, USA and Canada, while we plan to further extend this economy-linked data cloud in the near future.

Acknowledgements

This work has been supported by a national co-funded project (LinkedEconomy.org) under the 4th National Strategic Reference Framework (NSRF) period between 2007 and 2013. It was implemented by the Ministry of Education and Religious Affairs and the University of Piraeus Research Centre. The project aimed at providing a universal access to Greek and International economy data, as well as at promoting the benefits of linking heterogeneous sources under the concept of open and reusable data in this critical domain.

References

1. Michalis Vafopoulos, The Web economy: goods, users, models and policies, *Foundations and Trends in Web Science* 3, 1-2: 1–136, 2011.
2. M. Vafopoulos, M. Meimaris, Ioannis Anagnostopoulos, et al. 2015. Public spending as LOD: The case of Greece. *Semantic Web* 6, 2: 155–164.
3. Jakub; Klimek, Tomáš Knap, Jindrich Mynarz, Martin Necaský, and Vojtech Svátek. 2012. LOD2 - Creating Knowledge out of Interlinked Data Framework for Creating Linked Data in the Domain of Public Sector Contracts.
4. Seán O'Riain, Edward Curry, Andreas Harth, XBRL and open data for global financial ecosystems: A linked data approach, *International Journal of Accounting Information Systems*, Volume 13, Issue 2, June 2012, Pages 141-162
5. Hodess, Robin, Open Budgets: The Political Economy of Transparency, Participation, and Accountability, *Journal of Economic Literature* 52.2 (2014): 545-548.
6. Tygel AF, Attard J, Orlandi F, Campos ML, Auer S. "How much?" Is Not Enough-An Analysis of Open Budget Initiatives. arXiv preprint arXiv:1504.01563. 2015 Apr 7
7. Petrou I, Meimaris M, Papastefanatos G. Towards a methodology for publishing Linked Open Statistical Data. *JeDEM-eJournal of eDemocracy and Open Government*. 2014 Nov 14;6(1):97-105.
8. Höffner K, Martin M, Lehmann J. LinkedSpending: OpenSpending becomes Linked Open Data. *Semantic Web*. 2015 Mar 17;7(1):95-104.

9. Alvarez-Rodríguez, Jose María, Michalis Vafopoulos, and Juan Llorens, Enabling policy making processes by unifying and reconciling corporate names in public procurement data. The CORFU technique, Computer Standards & Interfaces 41 (2015): 28-38.
10. Vafopoulos MN, Vafeiadis G, Razis G, Anagnostopoulos I, Negkas D, Galanos L., Linked Open Economy: Take Full Advantage of Economic Data, available at SSRN 2732218. 2016 Feb 14.

Appendix

Query 1: Retrieving information from three different Greek data sources (Diavgeia – e-Procurement, and NSRF)

PREFIX elod: <<http://linkedeconomy.org/ontology#>>

PREFIX dcterms: <<http://purl.org/dc/terms/>>

PREFIX gr: <<http://purl.org/goodrelations/v1#>>

```

SELECT (COUNT(?spendingD) AS ?paymentDiavgeiaCounter)
(SUM(xsd:decimal(?amountD)) AS ?paymentDiavgeiaAmount)
(COUNT(?spendingK) AS ?paymentKhmdhsCounter) (SUM(xsd:decimal(?amountK))
AS ?paymentKhmdhsAmount) (COUNT(?subsidy) AS ?subsidyCounter)
(SUM(xsd:decimal(?siAmount)) AS ?subsidyAmount)
FROM <http://linkedeconomy.org/DiavgeiaII/2015>
FROM <http://linkedeconomy.org/EprocurementProper>
FROM <http://linkedeconomy.org/Subsidies>
WHERE {
#Transparency
{
?spendingD elod:hasExpenditureLine ?expLine ;
dcterms:issued ?date ;
rdf:type elod:SpendingItem .
?expLine elod:amount ?ups ;
elod:seller <http://linkedeconomy.org/resource/Organization/094019245> .
?ups gr:hasCurrencyValue ?amountD .
FILTER (?date >= "2015-01-01T00:00:00Z"^^xsd:dateTime) .
FILTER (?date < "2016-01-01T00:00:00Z"^^xsd:dateTime) .
FILTER NOT EXISTS {?spendingD elod:hasCorrectedDecision ?correctedDecision}
.
FILTER NOT EXISTS {?ups elod:riskError "1"^^xsd:boolean} .
}
UNION
#KHMDHS
{
?spendingK elod:hasRelatedContract ?contract ;
elod:hasExpenditureLine ?expLine .
?expLine elod:amount ?ups ;
elod:seller <http://linkedeconomy.org/resource/Organization/094019245> .

```

```

?ups gr:hasCurrencyValue ?amountK .
?contract elod:signatureDate ?date .
FILTER (?date >= "2014-01-01T00:00:00Z"^^xsd:dateTime) .
FILTER (?date < "2015-01-01T00:00:00Z"^^xsd:dateTime) .
}
UNION
#NSRF
{
?subsidy elod:hasRelatedSpendingItem ?spendingItem ;
  elod:beneficiary
<http://linkedeconomy.org/resource/Organization/094019245> .
?spendingItem elod:amount ?siUps .
?siUps gr:hasCurrencyValue ?siAmount.
}
}

```

Query 2: Combining eLod and Linked Open Data cloud information (from DBpedia)

```

PREFIX elod: <http://linkedeconomy.org/ontology#>
PREFIX gr: <http://purl.org/goodrelations/v1#>
PREFIX dbo: <http://dbpedia.org/ontology/>
PREFIX dcterms: <http://purl.org/dc/terms/>

SELECT (COUNT(?expenseApproval) AS ?expenseApprovalCounter)
(SUM(xsd:decimal(?amount)) AS ?expenseApprovalAmount) ?foundingYear
?netIncome ?numOfEmployees
WHERE
{
#Transparency
GRAPH <http://linkedeconomy.org/DiavgeiaII/2015> {
?expenseApproval elod:hasExpenditureLine ?expLine ;
  dcterms:issued ?date ;
  rdf:type elod:ExpenseApprovalItem .
?expLine elod:amount ?ups ;
  elod:seller <http://linkedeconomy.org/resource/Organization/094019245> .
?ups gr:hasCurrencyValue ?amount .
FILTER (?date >= "2015-01-01T00:00:00Z"^^xsd:dateTime) .
FILTER (?date < "2016-01-01T00:00:00Z"^^xsd:dateTime) .
FILTER NOT EXISTS {?expenseApproval elod:hasCorrectedDecision
?correctedDecision} .
}
#DBpedia
SERVICE <http://dbpedia.org/sparql> {
<http://dbpedia.org/resource/OTE> dbo:foundingYear ?foundingYear ;
  dbo:netIncome ?netIncome ; dbo:numberOfEmployees ?numOfEmployees .
}
}
}

```