1. Record Nr. UNINA9910484120503321

Titolo	The Semantic Web : 14th International Conference, ESWC 2017, Portorož, Slovenia, May 28 – June 1, 2017, Proceedings, Part I / / edited by Eva Blomqvist, Diana Maynard, Aldo Gangemi, Rinke Hoekstra, Pascal Hitzler, Olaf Hartig
Pubbl/distr/stampa	Cham : , : Springer International Publishing : , : Imprint : Springer, , 2017
ISBN	3-319-58068-X
Edizione	[1st ed. 2017.]
Descrizione fisica	1 online resource (XXX, 673 p. 151 illus.)
Collana	Information Systems and Applications, incl. Internet/Web, and HCI ; ; 10249
Disciplina	025.04
Soggetti	Information storage and retrieval Artificial intelligence Application software Computers Data mining Software engineering Information Storage and Retrieval Artificial Intelligence Information Systems Applications (incl. Internet) Theory of Computation Data Mining and Knowledge Discovery Software Engineering
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Intro -- Preface -- Organization -- Abstract of Keynotes -- Bringing Semantic Intelligence to Financial Markets -- Disrupting the Semantic Comfort Zone -- Semantic Web Technologies for Digital Archives -- Contents - Part I -- Contents -- Part II -- Semantic Data Management, Big Data, and Scalability Track -- Traffic Analytics for Linked Data Publishers -- 1 Introduction -- 2 Related Work -- 3 System Overview -- 4 Traffic Metrics -- 4.1 Content Metrics Extraction -- 4.2 Protocol Metrics Extraction and SPARQL Queries Weight -- 4.3 Audience Metrics

Extraction and Visitor Session Identification -- 5 Results -- 6 Conclusions and Future Work -- References -- Explaining Graph Navigational Queries -- 1 Introduction -- 2 Related Work -- 3 Building Query Explanations with GeL -- 3.1 Syntax of GeL -- 3.2 Semantics of GeL -- 4 Algorithms and Complexity -- 4.1 Translating GeL into SPARQL -- 5 Implementation and Evaluation -- 6 Concluding Remarks and Future Work -- References -- A SPARQL Extension for Generating RDF from Heterogeneous Formats -- 1 Introduction -- 2 Use-Cases and Requirements -- 3 Related Work -- 4 SPARQL-Generate Specification -- 4.1 SPARQL-Generate Concrete Syntax -- 4.2 Abstract Syntax -- 4.3 SPARQL-Generate Semantics -- 5 Implementation and Evaluation -- 5.1 Generic Approach -- 5.2 Implementation on Top of Apache Jena -- 5.3 Evaluation -- 6 Conclusion and Future Work -- References -- Linked Data Track -- Exploiting Source-Object Networks to Resolve Object Conflicts in Linked Data -- 1 Introduction -- 2 Preliminaries -- 2.1 Basic Definitions -- 2.2 Problem Analysis -- 3 ObResolution Method -- 3.1 Model Details -- 3.2 Inference Algorithms -- 3.3 Practical Issues -- 4 Evaluation -- 4.1 The Datasets -- 4.2 Comparative Methods and Metrics -- 4.3 Results -- 5 Related Work -- 6 Conclusion and Future Work -- References.
Methods for Intrinsic Evaluation of Links in the Web of Data -- 1 Introduction -- 2 Preliminaries -- 3 Principles for Data Interlinking in the Web of Data -- 4 Intrinsic Measures for Assessing the Quality of Links -- 4.1 Basic Descriptive Statistics -- 4.2 Principles-Based Measures -- 5 Empirical Analysis -- 5.1 Data -- 5.2 Methodology -- 5.3 Measure Validation -- 5.4 Results -- 6 Related Work -- 7 Conclusions and Future Work -- References -- Entity Deduplication on ScholarlyData -- 1 Introduction -- 2 Related Work -- 3 Deduplication -- 3.1 Blocking Strategies -- 3.2 Classification -- 3.3 URI Harmonisation -- 4 Experiments -- 4.1 The Train/test Dataset -- 4.2 Blocking -- 4.3 Classification -- 4.4 URI Harmonisation -- 5 Conclusions -- References -- Machine Learning Track -- WOMBAT -- A Generalization Approach for Automatic Link Discovery -- 1 Introduction -- 2 Preliminaries -- 3 Constructing and Traversing Link Specifications -- 3.1 Learning Atomic Specifications -- 3.2 Combining Atomic Specifications -- 4 The WOMBAT Algorithm -- 5 Evaluation -- 6 Related Work -- 7 Conclusions and Future Work -- References -- Actively Learning to Rank Semantic Associations for Personalized Contextual Exploration of Knowledge Graphs -- 1 Introduction -- 2 Contextual KG Exploration -- 3 Active Learning to Rank for Semantic Associations -- 3.1 Features -- 4 Experiments -- 4.1 Experimental Settings -- 4.2 Configurations and Baselines -- 4.3 Results and Discussion -- 5 Related Work -- 6 Conclusion -- References -- Synthesizing Knowledge Graphs for Link and Type Prediction Benchmarking -- 1 Introduction -- 2 Related Work -- 3 Knowledge Graph Model -- 4 Synthesis Process -- 5 Experiments -- 6 Conclusion and Outlook -- References -- Online Relation Alignment for Linked Datasets -- 1 Introduction -- 2 Related Work -- 3 Preliminaries -- 4 Problem Statement.
5 SORAL: Relation Alignment -- 5.1 Candidate Generation -- 5.2 Sampling Strategies -- 5.3 Features -- 6 Experimental Setup -- 7 Results and Discussion -- 7.1 Relation Alignment Accuracy -- 7.2 Query-Execution Overhead -- 8 Conclusion -- References -- Tuning Personalized PageRank for Semantics-Aware Recommendations Based on Linked Open Data -- 1 Introduction -- 2 Related Work -- 3 Methodology -- 3.1 Graph-Based Representations -- 3.2 Running Personalized PageRank -- 4 Experimental Evaluation -- 4.1 Experimental Protocol -- 4.2 Discussion of the Results -- 5

Conclusions and Future Work -- References -- Terminological Cluster Trees for Disjointness Axiom Discovery -- 1 Introduction -- 2 Related Work -- 3 Disjointness Discovery as a Conceptual Clustering Problem -- 4 Terminological Cluster Trees for Disjointness Learning -- 4.1 Growing Terminological Cluster Trees -- 4.2 Extracting Candidate Disjointness Axioms from TCTs -- 5 Experiments -- 5.1 Re-discovery of a Target Disjoint Axiom -- 5.2 Comparison to Other Approaches Under SDA -- 6 Conclusions and Outlook -- References -- Embedding Learning for Declarative Memories -- 1 Introduction -- 2 Unique-Representation Hypothesis -- 3 Semantic and Episodic Knowledge Graph Models -- 3.1 Semantic Knowledge Graph -- 3.2 An Event Model for Episodic Memory -- 4 Tensor Decompositions -- 4.1 Tensor Decompositions -- 4.2 Inner Product Formulation of Tensor Decompositions -- 5 Querying Memories -- 5.1 Probabilistic Querying -- 5.2 Semantic Memory Derived from Episodic Memory -- 6 Relationships to Human Memories -- 6.1 Unique-Representation Hypothesis for Entities and Predicates -- 6.2 Perception and Memory Formation -- 6.3 Tensor Memory Hypothesis -- 6.4 Semantic Memory and Episodic Memory -- 7 Experiments -- 7.1 Data Set -- 7.2 Evaluation and Implementation -- 7.3 Experimental Results -- 8 Conclusions -- References.
Mobile Web, Sensors, and Semantic Streams Track -- Spatial Ontology-Mediated Query Answering over Mobility Streams -- 1 Introduction -- 2 V2X Integration using a Local Dynamic Map -- 3 Streams, Pulses, and Spatial Databases -- 4 Syntax, Semantics, and Query Language of DL-LiteA (S,F) -- 5 Query Rewriting by Stream Aggregation -- 6 Query Evaluation by Hypertree Decomposition -- 7 Implementation and Experimental Evaluation -- 8 Related Work and Conclusion -- References -- Optimizing the Performance of Concurrent RDF Stream Processing Queries -- 1 Introduction -- 2 Foundations -- 2.1 Multi-way Join -- 2.2 Shared Join Operator and Network of Shared Join Operators -- 3 Optimization for Concurrent CQELS Queries -- 3.1 CQELS+: Network of Shared Join Operators -- 3.2 Load Balancing for Parallel CQELS+ Instances -- 4 Evaluation -- 4.1 Evaluating Shared Joins in CQELS+ -- 4.2 Evaluating Load Balancing over CQELS+ -- 4.3 Evaluating the Query Registration Time -- 5 Related Work -- 6 Conclusions and Future Work -- References -- AGACY Monitoring: A Hybrid Model for Activity Recognition and Uncertainty Handling -- 1 Introduction -- 2 Related Work -- 3 The AGACY Monitoring Architecture Overview -- 4 Knowledge Based Layer -- 4.1 Ontological Modeling -- 4.2 Semantic Reasoning -- 5 Data Driven Layer -- 5.1 Time and Uncertainty-Based Features Extraction -- 5.2 Dempster-Shafer Theory for Activity Classification -- 5.3 Activities Instances Inferring Under Uncertainty -- 6 Evaluation and Discussion -- 6.1 DataSet -- 6.2 Implementation and Experimental Setup -- 6.3 Evaluation and Results -- 7 Conclusion and Future Work -- References -- Natural Language Processing and Information Retrieval Track -- Mapping Natural Language to Description Logic -- 1 Introduction -- 2 Related Work -- 3 Approach Overview -- 3.1 Grammar -- 3.2 Semantic Parser and Surface Realiser.
4 Evaluation and Results -- 4.1 Mapping SIDPs to Complex Axioms -- 4.2 Assessing Correctness -- 4.3 Ontology Enrichment -- 5 Conclusion -- References -- Harnessing Diversity in Crowds and Machines for Better NER Performance -- 1 Introduction -- 2 Use Case and Datasets -- 3 Related Work -- 3.1 Open Knowledge Extraction Systems -- 3.2 Crowdsourcing Named Entities -- 3.3 Multi-NER, Hybrid Named Entity Recognition -- 4 Single-NER vs. Multi-NER Comparison -- 4.1 Single-NER vs. Multi-NER - Entity Surface -- 4.2 Single-NER vs. Multi-NER -

Entity Surface and Entity Type -- 4.3 Analysis of False Negative Named Entities -- 4.4 Analysis of False Positive Named Entities -- 5 Experimental Setup -- 5.1 Crowdsourcing Experimental Data -- 5.2 Crowdsourcing Annotation Task -- 5.3 CrowdTruth Metrics -- 6 Results -- 7 Discussion -- 8 Conclusion -- References -- All that Glitters Is Not Gold -- Rule-Based Curation of Reference Datasets for Named Entity Recognition and Entity Linking -- 1 Introduction -- 2 Related Work -- 3 Formal Annotation Framework -- 3.1 Assumptions -- 3.2 Rule Set -- 3.3 Comparison with Related Work -- 3.4 Observations -- 4 Eaglet -- 4.1 Preprocessing Module -- 4.2 Completion Module -- 4.3 Error Detection Pipeline -- 4.4 Review Module -- 5 Evaluation -- 5.1 Experiment I -- 5.2 Experiment II -- 5.3 Experiment III -- 6 Conclusion -- References -- Semantic Annotation of Data Processing Pipelines in Scientific Publications -- 1 Introduction -- 2 Related Work -- 3 The DMS Ontology -- 4 DPP Knowledge Extraction Workflow -- 4.1 Training Data Generation -- 4.2 Classification and NER -- 4.3 Linked Data Generation -- 5 Evaluation -- 5.1 Dataset -- 5.2 Analysis of Rhetorical Classifiers -- 5.3 Quality of Extracted Entities -- 6 Conclusion -- References -- Combining Word and Entity Embeddings for Entity Linking -- 1 Introduction -- 2 Related Work.
3 Combining Word and Entity Embeddings.

Sommario/riassunto

The two volumes LNCS 10249 and 10250 constitute the refereed proceedings of the 14th International Semantic Web Conference, ESWC 2017, held in Portorož, Slovenia. The 51 revised full papers presented were carefully reviewed and selected from 183 submissions. In addition, 10 PhD papers are included, selected out of 14 submissions. The papers are organized in the following tracks: semantic data management, big data, and scalability; linked data; machine learning; mobile web, sensors, and semantic streams; natural language processing and information retrieval; vocabularies, schemas, and ontologies; reasoning; social web and web science; semantic web and transparency; in use and industrial track; and PhD symposium. The paper 'Linked Data Notifications: A Resource-Centric Communication Protocol' is published open access under a CC BY 4.0 license at link. springer.com. .