Chapter 11

Computational Cognitive Systems

Timo Honkela, Krista Lagus, Marcus Dobrinkat, Oskar Kohonen, Mikaela Kumlander, Tiina Lindh-Knuutila, Ilari Nieminen, Mari-Sanna Paukkeri, Matti Pöllä, Juha Rautio, Sami Virpioja, Jaakko Väyrynen, Paul Wagner, Eric Malmi, Tero Tapiovaara, Tommi Vatanen, Ilkka Kivimäki, Laura Leppänen, Sini Pessala, Santosh Tirunagari

11.1 Introduction

Computational Cognitive Systems group conducts research on artificial systems that combine perception, action, reasoning, learning and communication. This area of research draws upon biological, cognitive and social system approaches to understanding cognition. Cognitive systems research is multidisciplinary and interdisciplinary. It benefits from sharing and leveraging expertise and resources between disciplines. Methodologically, statistical machine learning, pattern recognition and signal processing are central tools within computational cognitive systems research. Our research focuses on modeling and applying methods of unsupervised and semisupervised learning in the areas of conceptual modeling, multilingual and language independent language technology, and socio-cognitive modeling. Results related to language processing are reported in Section 10.

We approach conceptual modeling as a dynamic phenomenon. Among humans, conceptual processing takes place as an individual and social process. We attempt to model this dynamic and constructive aspect of conceptual modeling by using statistical machine learning methods. We also wish to respect the overall complexity of the theme, for instance, not relying on explicit symbolic representations are the only means relevant in conceptual modeling. Our machine translation research builds on the conceptual modeling research as well as on the research on adaptive language technology.

Socio-cognitive modeling is our newest research area which builds on 1) the experience and expertise in modeling complex phenomena related to language learning and use at cognitive and social levels and 2) strong national and international collaboration especially with the representatives of social sciences and humanities. Socio-cognitive modeling mainly merges aspects of computer science, social sciences and cognitive science. The basic idea is to model interlinked social and cognitive phenomena.

Summary of collaboration

We have worked in close collaboration with other groups in Adaptive Informatics Research Centre, lead by Prof. Erkki Oja and Prof. Samuel Kaski, in particular natural language processing and multimodal interfaces (Dr. Mikko Kurimo and Dr. Jorma Laaksonen).

The collaboration with *Aalto School of Economics* and *National Consumer Research Centre* that started in KULTA projects, has continued within Tekes-funded VirtualCoach project. The project focuses on wellbeing informatics and is discussed below in more detail.

In the area of multilingual language technology, META-NET Network of Excellence is a major effor (http://www.meta-net.eu). One objective is to build bridges to neighbouring technology fields such as machine learning and cognitive systems. The research agenda consist of four areas: (1) bringing more semantics into Machine Translation, (2) optimising the division of labour in hybrid Machine Translation, (3) exploiting the context for Translation, and (4) preparing a base for Machine Translation. The COG group has been actively involved in the third area. In overall, META-NET consists of 57 research centres from 33 countries. META-NET is coordinated by the German Research Center for Artificial Intelligence (DFKI). A Cognitive Systems blog description of a META-NET event, written by Jaakko Väyrynen is shown in Fig. 11.1 (http://cogsys.blogspot.com/2011/06/meta-forum-2011.html).

META-FORUM 2011

The META-FORUM 2011 is held June 27/28 in Budapest, Hungary. It is organized by META-NET, a Network of Excellence consisting of 47 research centers from 31 countries, which is dedicated to building the technological foundations of a multilingual European information society. The forum consists of keynote talks, several sessions, two panels and joined exhibition by industry and academia.

Thomas Hofmann from Google Research gave a keynote lecture on how they are bringing semantics to the web. He explained how interesting models can be scaled to web size by using curated knowledge bases (eg Freebase) and the output from accurate but slower parsers as the either the starting point to statistical analysis or training data for faster systems.

Bran Boguraev from IBM's TJ Watson Research Center talked about evidence based reasoning over natural language content. The example was of course their massively parallel Watson architecture for large-scale hypothesis generation, validation and scoring which tackles the long-standing challenge in AI to emulate human expertise.

Hans Uszkoreit introduced the outline for META-NET's strategic research agenda for language technology in Europe. The vision paper presents possible futures for innovative language technology applications in the following areas:

- language-transparent web and media experience,
- natural and inclusive interaction, and
- efficient information management.

Posted by Jaakko Väyrynen at 3:37 PM

Figure 11.1: META-FORUM 2011 description in the widely read Cognitive Systems blog.

The COG group has actively participated the MultilingualWeb initiative that is concerned with standards and best practices that support the creation, localization and use of multilingual web-based information. The consortium is coordinated by W3C and includes companies such as Microsoft, Facebook, Opera and SAP. Fig. 11.2 gives an excerpt of a blog post to Cognitive Systems on a MultilingualWeb event, written by Matti Pöllaä (see http://cogsys.blogspot.com/2011/04/content-on-multilingual-web.html for more details).

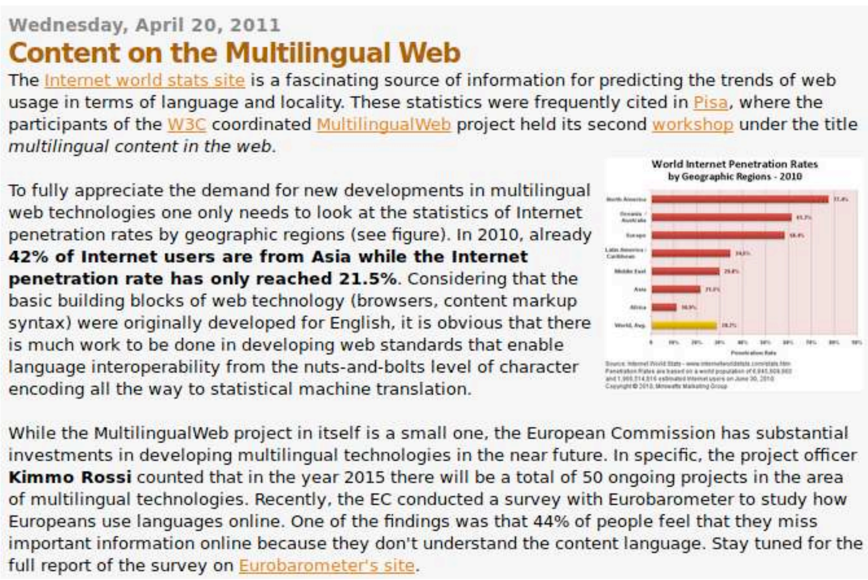

Wednesday, April 20, 2011

Content on the Multilingual Web

The Internet world stats site is a fascinating source of information for predicting the trends of web usage in terms of language and locality. These statistics were frequently cited in Pisa, where the participants of the W3C coordinated MultilingualWeb project held its second workshop under the title *multilingual content in the web*.

To fully appreciate the demand for new developments in multilingual web technologies one only needs to look at the statistics of Internet penetration rates by geographic regions (see figure). In 2010, already **42% of Internet users are from Asia while the Internet penetration rate has only reached 21.5%**. Considering that the basic building blocks of web technology (browsers, content markup syntax) were originally developed for English, it is obvious that there is much work to be done in developing web standards that enable language interoperability from the nuts-and-bolts level of character encoding all the way to statistical machine translation.

While the MultilingualWeb project in itself is a small one, the European Commission has substantial investments in developing multilingual technologies in the near future. In specific, the project officer **Kimmo Rossi** counted that in the year 2015 there will be a total of 50 ongoing projects in the area of multilingual technologies. Recently, the EC conducted a survey with Eurobarometer to study how Europeans use languages online. One of the findings was that 44% of people feel that they miss important information online because they don't understand the content language. Stay tuned for the full report of the survey on Eurobarometer's site.

Figure 11.2: A description of a MultilingualWeb event.

MultilingualWeb has been organizing workshops open to the public and various communication channels, spreading information about what standards and best practices currently exist, and what gaps need to be filled. Fig. 11.3 shows a fragment of the result a thematic session. Its focus was "Semantic resources and machine learning for quality, efficiency and personalisation of accessing relevant information over language borders".

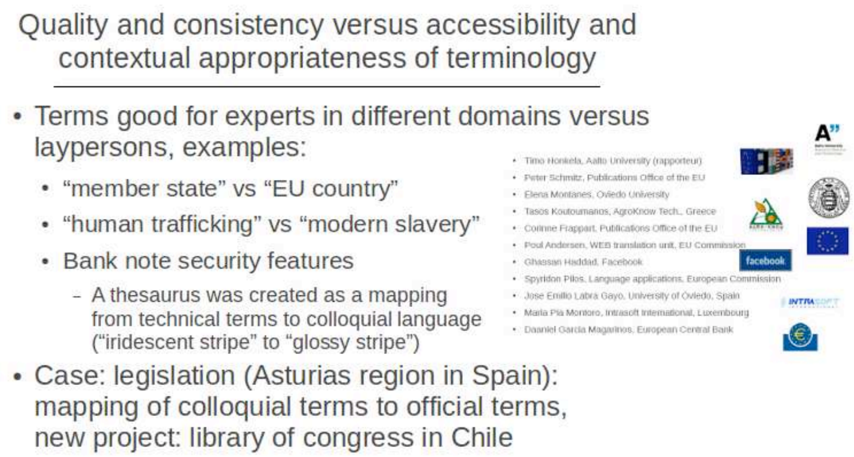

Figure 11.3: A fragment of a result from a MultilingualWeb thematic session.

An excellent example of the positive effects of international researcher exchange are the results that stem from the visit by young COG researchers, Tommi Vatanen and Eric Malmi, to CERN in Switzerland [5, 5].

In EIT ICT Labs, Dr. Krista Lagus has served as the Lead of Schools & Camps Catalyst Development. EIT ICT Labs is one of the first three Knowledge and Innovation Communities (KICs) selected by the European Institute of Innovation & Technology (EIT) to accelerate innovation in Europe. EIT aims to rapidly emerge as a key driver of EU's sustainable growth and competitiveness through the stimulation of world-leading innovation (http://eit.ictlabs.eu/).

Scientific events

In collaboration with other groups in AIRC, the COG group has been active in organizing national and international conferences. Two main events took place in 2011, International Conference on Artificial Neural Networks [49, 50] and Workshop on Self-Organizing Maps [3].

During ICANN 2011, META-NET workshop on Context in Machine Translation was organized to foster exchange of ideas and results in this area. The notion of context was meant to be understood broadly, including other modalities (like vision) in addition to the textual contexts.

The Context in Machine Translation Challenge is part of a series of challenges organized by the META-NET Network of Excellence (http://www.meta-net.eu), jointly by Aalto University (Finland), CNRS/LIMSI (France) and ILSP (Greece), supported by other network partners.

The COG group was also involved in organizing the Finnish Artificial Intelligence Conference, STeP 2011 [4].

VirtualCoach project

VirtualCoach – Paths of Wellbeing is a major project in the area of wellbeing informatics, lead by Dr. Krista Lagus (http://blog.pathsofwellbeing.com/). The VirtualCoach builds on the traditional methodological strengths of the COG group and AIRC, in general. Wellbeing informatics is an emerging area of research in which ICT methodologies are used to measure, analyze, and promote wellbeing of individuals. Examples of traditional applications include heart rate monitoring, tracking sports activities, analyzing the nutritional content of diets, and analyzing sleeping patterns with mobile technologies.

In the VirtualCoach project, a central topic is how to help people to find peer and professional support in a personalized manner. One approach is to build social media applications in which users can find stories that are potentially helpful in their individual life situations. The users may wish to develop their wellbeing further, or need to solve some problem that prevents them from achieving a satisfactory level of wellbeing.

The VirtualCoach project is a collaborative research effort with the National Consumer Research Center and several companies including createAmove, FlowDrinks, Futuria Consulting, If insurance company, Innotiimi, mutual pension insurance company Varma, MTV Media, Oppifi, Terveystalo Healthcare, and Vierumäki Sports Institute.

References

[1] Timo Honkela, Wlodzislaw Duch, Mark A. Girolami, and Samuel Kaski, editors (2011). Artificial Neural Networks and Machine Learning - Proceedings of ICANN 2011 - 21st International Conference on Artificial Neural Networks, Parts I and II. Springer.

[2] Mikael Kuusela, Eric Malmi, Risto Orava, and Tommi Vatanen (2011). Soft classification of diffractive interactions at the LHC. AIP Conference Proceedings, 1350(1):111-114.

[3] Jorma Laaksonen and Timo Honkela, editors (2011). Advances in Self-Organizing Maps - Proceedings of WSOM 2011, 8th International Workshop. Springer.

[4] Tapio Pahikkala, Jaakko Väyrynen, Jukka Kortela, and Antti Airola, editors (2010). Proceedings of the 14th Finnish Artificial Intelligence Conference, STeP 2010, Finnish Artificial Intelligence Society.

[5] Tommi Vatanen, Mikael Kuusela, Eric Malmi, Tapani Raiko, Timo Aaltonen, and Yoshikazu Nagai (2011). Fixed-background EM algorithm for semi-supervised anomaly detection. Technical report, Aalto University School of Science.

11.2 Learning to translate

Our research on multilinguality and machine translation (MT) uses novel methods that are based on adaptivity. An MT system is *learning to translate* rather than needs to be programmed to do so. The advances in statistical machine translation have shown that the adaptive paradigm can help in reducing the system development costs dramatically. However, these systems rely on representations that do not capture many relevant linguistic aspects, neither take into account the wealth of knowledge that is known about human cognitive processes related to natural language understanding, translation and interpretation.

An important context for the research and development work on multilinguality is META-NET. META-NET, a Network of Excellence consisting of 57 research centres from 33 countries, is dedicated to building the technological foundations of a multilingual European information society [1]. The research work in META-NET has been divided into four work packages. Our activities have focused on exploiting context in machine translation. During ICANN 2011 conference, a META-NET workshop on "Context in Machine Translation" was organized[2]. The objective was to foster exchange of ideas and results in this area. Here the notion of context is meant to be understood broadly, including other modalities (like vision) in addition to the textual contexts. An invited talk was given by Dr. Katerina Pastra entitled "Bridging language, action and perception: the cognitive context of machine translation". During the workshop a challenge on context in mt was announced. The challenge data set consists of documents from the JRC-ACQUIS Multilingual Parallel Corpus. Two language-pair directions are included in the data, English-¿Finnish and Greek-French. The constructed challenge training data set contains the document context, n-best lists for translated documents and additional contextual information as well as the reference translations.

Language identification of short text segments

For processing multilingual texts, it is important to identify the language of each document, sentence, or even word. There are many accurate methods for language identification of long text samples, but identification of very short strings still presents a challenge. In [1], we consider test samples that have only 5–21 characters. We show that a simple but efficient method, naive Bayes classifier based on character n-gram models, outperforms previous methods, when state-of-the-art language modeling techniques from automatic speech recognition research are applied. Using the Universal Declaration of Human Rights as a data set, we were able to conduct the experiments with as many as 281 languages.

Automatic machine translation evaluation

The normalized compression distance (NCD) has been further investigated as an automatic machine translation metric. It is based on character-sequence comparison of translated text and a reference translation, whereas most typical metrics (e.g. BLEU and NIST) operate on word-sequences. In [2], the NCD metric has been extended to include flexible word matching, which extends the references translations with synonyms for words.

[1] http://www.meta-net.eu/

[2] http://www.cis.hut.fi/icann11/con-txt-mt11/

Several possible extensions were tested in [3] in order to improve the evaluation metric, including multiple reference handling and segment replication. The metric also participated in the MetricsMATR 2010 machine translation evaluation shared task. MT evaluation metrics are themselves evaluated by measuring correlation between the automatic metric and known human evaluations of translations.

Automatic evaluation of machine translation (MT) systems requires automated procedures to ensure consistency and efficient handling of large amounts of data, and are essential for parameter optimization and statistical machine translation system development. In contrast to most MT evaluation measures, e.g., BLEU and METEOR, NCD provides a general information theoretic measure of string similarity.

NCD is an approximation of the uncomputable normalized information distance (NID), a general measure for the similarity of two objects. NID is based on the notion of Kolmogorov complexity, a theoretical measure for the information content of a string, defined as the shortest universal Turing machine that prints the string and stops. NCD approximates NID by the use of a compressor that is an upper bound of the Kolmogorov complexity.

Similar to the mBLEU extension of the BLEU metric, the same synonym handling module from METEOR was incorporated into the NCD metric. In our experiments, the resulting mNCD metric had consistently better correlation with human judgments of translation compared to the basic NCD metric.

An NCD-based metric was developed to handle multiple references in evaluation. It can be viewed as a generalization of the NCD metric, as they are equal with only one reference translations. It was shown to work better when two reference translations are available.

Domain adaptation for statistical machine translation

Statistical machine translation methodology is highly dependent of relevant parallel texts for training. However, available large parallel corpora are typically out-of-domain for many interesting translation tasks, such as news translation. This is especially true for less-resourced languages. Therefore methods that can utilize out-of-domain text

Four existing different domain adaptation methods were tested in [4]: language model (LM) adaptation, translation model (TM) adaptation, automatic post-editing and retraining with combined data. All tested methods except language model adaptation outperformed the baseline system trained with only the out-of domain data. The experiment were conducted with a larger out-of-domain Europarl parallel corpus and a small previously collected small corpus of Finnish Iltalehti news with their English translations.

Domain adaption can be accomplished at several locations in the statistical machine translation process. The simplest way is simply to pool all available data and to learn a single model based on it. It may not feasible if only models are available or the models are incompatible. Also, it may give too little emphasis on the small amount of in-domain data. *Language model adaptation* requires only monolingual target language data and affects only the selection of translations without providing any new translations possibilities for words or phrases. The adaptation can be done with a combination of the data or linear or log-linear interpolation of two or more language models. *Translation model adaptation* requires additional parallel data that can be either included in the existing data or the models can be joined with log-linear interpolation. The process is illustrated in

Figure 11.4. The *post-edit domain-adaption*, shown in Figure 11.5, learns another translation model from the output of the original translation system to correct or corrected translations, trying to statistically fix mistakes made by the original system.

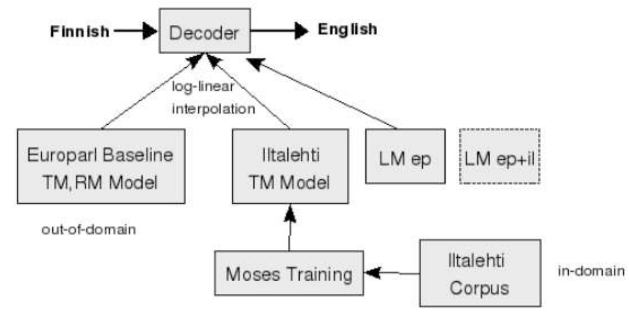

Figure 11.4: The process for domain adaptation with log-linear interpolation of baseline and in-domain translation models (TM). The wwo models operate on parallel inside one translation system.

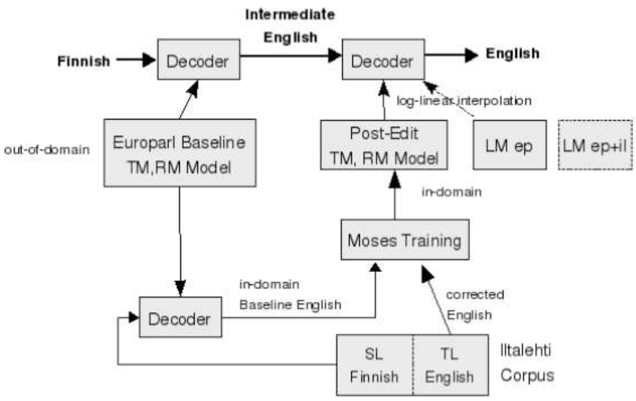

Figure 11.5: The process for post-edit domain adaptation. The two translation systems operate in sequence.

References

[1] Tommi Vatanen, Jaakko J. Väyrynen, and Sami Virpioja. Language identification of short text segments with n-gram models. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Bente Maegaard, Joseph Mariani, Jan Odjik, Stelios Piperidis, Mike Rosner, and Daniel Tapias, editors, *Proceedings of the Seventh conference on International Language Resources and Evaluation (LREC'10)*, Valletta, Malta, May 2010. European Language Resources Association (ELRA).

[2] Marcus Dobrinkat, Tero Tapiovaara, Jaakko Väyrynen, and Kimmo Kettunen. Evaluating machine translations using mNCD. In *Proceedings of the ACL 2010 Conference Short Papers*, pages 80–85. Association for Computational Linguistics, 2010.

[3] Marcus Dobrinkat, Tero Tapiovaara, Jaakko Väyrynen, and Kimmo Kettunen. Normalized compression distance based measures for MetricsMATR 2010. In *Proceedings*

of the Joint Fifth Workshop on Statistical Machine Translation and MetricsMATR, pages 343–348. Association for Computational Linguistics, 2010.

[4] Marcus Dobrinkat and Jaakko J. Väyrynen. Experiments with domain adaptation methods for statistical MT: From european parliament proceedings to finnish newspaper text. In *Proceedings of the 14th Finnish Artificial Intelligence Conference STeP 2010*, number 25 in Publications of the Finnish Artificial Intelligence Society, pages 31–38. Finnish Artificial Intelligence Society, 2010.

11.3 Socio-cognitive modeling

Socio-cognitive modeling is a new research area that merges aspects of computer science, social sciences and cognitive science. The basic idea is to model interlinked social and cognitive phenomena. Our focus has traditionally been in modeling individual cognition that learns and uses language, or in building models of language using statistical machine learning methods. Already for a long time, we have been interested in language and its use as a dynamic phenomenon rather than as a static structural object. Thereafter, we have widened our interest to language as a socio-cultural phenomenon that encodes human knowing and further to other socio-cognitive phenomena, however often related to language. In other words, cognition and intelligent activity are not only individual processes but ones which rely on socio-culturally developed cognitive tools. These include physical and conceptual artifacts as well as socially distributed and shared processes of intelligent activity embedded in complex social and cultural environments.

One approach in socio-cognitive modeling is social simulation. It aims at exploring and understanding of social processes by means of computer simulation. Social simulation methods can be used to to support the objective of building a bridge between the qualitative and descriptive approaches used in the social sciences and the quantitative and formal approaches used in the natural sciences. Collections of agents and their interactions are simulated as complex non-linear systems, which are difficult to study in closed form with classical mathematical equation-based models. Social simulation research builds on the distributed AI and multi-agent system research with a specific interest of linking the two areas.

Directions for e-science and science 2.0

Science 2.0 builds on the technologies of Web 2.0. Blogs, wikis and other social sharing and interaction tools allow scientists to interact and make their data and interpretations available for others in novel ways [1]. Science 2.0 continues and extends the tradition of publishing open source software and open access publishing of scientific articles. Local examples, stemming from the research of AIRC and its predecessors, include *SOM Toolbox* for Matlab (`http://www.cis.hut.fi/somtoolbox/`), *FastICA* for Matlab (`http://www.cis.hut.fi/projects/ica/fastica/`), *dredviz* software package for dimensionality reduction in information visualization (`http://www.cis.hut.fi/projects/mi/software/dredviz/`), and *Morfessor* software (`http://www.cis.hut.fi/projects/morpho/`).

Where Science 2.0 refers to new practices in conducting science with the help of communications and collaborations technologies, computational science builds on modeling and simulation of real world or anticipated phenomena based on massive data sets. Traditionally, this field has been dominated by applications related to natural sciences and engineering, but also human and social sciences have started to use computational models as a research tool.

An article, based on a keynote given by Timo Honkela in MASHS 2010 conference, discusses the themes described above and considers computational linguistics and computational economics as more specific examples [1]. Also a map of Finnish science was described and discussed (see Fig. 11.6). The map was created based on the contents of 3,224 application documents sent to Academy of Finland. The text contents were analyzed using an automatic terminology extraction method called Likey (see Section 10.1 and the section on keyphrase extraction for more detauls). The SOM algorithm organized the documents into a map in which similar applications are close to each other and in which thematic areas emerged (Fig. 11.6).

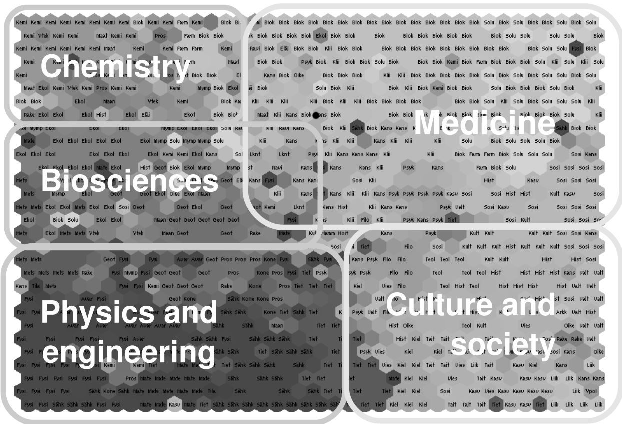

Figure 11.6: Map of Finnish science.

Text mining and qualitative analysis of history interviews

In collaboration with Dr. Petri Paju, we have explored the possibility of applying a text mining methods on a large qualitative source material concerning the history of information technology [3]. This data was collected in the Swedish documentation project "From Computing Machines to IT." We applied text mining on the interview transcripts of this Swedish documentation project. Specifically, we seeked to group the interviews according to their central themes and affinities and pinpoint the most relevant interviews for specific research questions. In addition, we searched for interpersonal links between the interviews. We applied the SOM algorithm to create a similarity diagram of the interviews. In the article based on this research, we discussed the results in several contexts including the possible future uses of text mining in researching history [3].

Analysis of global nutrition, lifestyle and health situation

EIT ICT Labs Wellbeing Innovation Camp (WIC) 2010 lead into a number of results (see http://www.cis.hut.fi/wicamp/2010/). In two cases, research work leading to publications in the general area of wellbeing informatics [2, 4]. The Action Line "Health and Wellbeing" in EIT ICT Labs has very similar basic objectives as the VirtualCoach project, discussed earlier in this report (see Section 11.1). In this action line, it is acknowledged that health and wellbeing "needs to be approached in a holistic way that fosters mental and physical fitness and balance. Having healthy and caring relationships, as well as good daily habits and behavioural patterns, are just two of the many principles in this holistic approach." (http://eit.ictlabs.eu/action-lines/health-wellbeing/)

In the first Wellbeing Innovation Camp project, the relationship between nutrition, lifestyle and health situation around the world was studied. The dataset used in the analysis is comprised of statistics that can be divided into three categories namely health, diet and lifestyle. The first category contains information such as obesity prevalence, incidence of tuberculosis, mortality rates and related variables in different countries. The dietary information includes consumption of proteins, sugar and milk products, and various other components of nutrition. The lifestyle category provides information related to the drinking and smoking habits, etc. One interesting finding was that there is a clear correlation between high consumption of sugar or sweeteners and a prevalence of cholesterol in men and women. This and other conclusions from the study are reported in [2].

Emotional semantics of abstract art and low-level image features

The second result of a project initiated at the Wellbeing Innovation Camp 2010, is an analysis on how well low-level image features can be used in predicting the emotional responses of subjects to abstract art [4]. This research was conducted in close collaboration with Jorma Laaksonen's research group. In this work, we studied people's emotions evoked by viewing abstract art images based on traditional low-level image features within a binary classification framework. Abstract art is used instead of artistic or photographic images because those contain contextual information that influences the emotional assessment in a highly individual manner. Whether an image of a cat or a mountain elicits a negative or positive response is subjective. After discussing challenges concerning image emotional semantics research, we empirically demonstrated that the emotions triggered by viewing abstract art images can be predicted with reasonable accuracy by machine using a variety of low-level image descriptors such as color, shape, and texture [4].

References

[1] Timo Honkela. Directions for e-science and science 2.0 in human and social sciences (2010). In Proceedings of MASHS 2010, Computational Methods for Modeling and Learning in Social and Human Sciences, pages 119–134. Multiprint.

[2] Yasir Mehmood, Mudassar Abbas, Xi Chen, and Timo Honkela (2011). Self-organizing maps of nutrition, lifestyle and health situation in the world. In Advances in Self-Organizing Maps - Proceedings of WSOM 2011, 8th International Workshop, pages 160–167. Springer.

[3] Petri Paju, Eric Malmi, and Timo Honkela (2011). Text mining and qualitative analysis of an it history interview collection. In History of Nordic Computing, IFIP Publications, pages 433–443. Springer.

[4] He Zhang, Eimontas Augilius, Timo Honkela, Jorma Laaksonen, Hannes Gamper, and Henok Alene (2011). Analyzing emotional semantics of abstract art using low-level image features. In Joao Gama, Elizabeth Bradley, and Jaakko Hollmén, editors, Advances in Intelligent Data Analysis X, volume 7014 of Lecture Notes in Computer Science, pages 413–423, Berlin/Heidelberg, Springer.

11.4 GICA: Grounded Intersubjective Concept Analysis

We have introduced a novel method to analyze and make visible differences among people regarding how they conceptualize the world [1, 3]. The Grounded Intersubjective Concept Analysis (GICA) method first employs either a conceptual survey or a text mining step to elicit particular ways in which terms and associated concepts are used among individuals. The subsequent analysis and visualization reveals potential underlying groupings of people, objects and contexts. The GICA method extends the basic idea of the traditional term-document matrix analysis to include a third dimension of different individuals. This leads to a formation of a third-order tensor of Subjects x Objects x Contexts. Through flattening, these Subject-Object-Context (SOC) tensors can be analyzed using various computational methods. In the following, we introduce the GICA method and its background in some detail.

Introduction to epistemological subjectivity

When human communication as well as computational modeling of knowledge and language is considered, it is usually taken granted that the meaning of all symbols used in the communication or representation of knowledge is shared by all human and/or artificial agents. It is, however, quite straightforward to show empirically that this is not the case. Practical and theoretical limitations of traditional knowledge representation were already highlighted in an early project the objective of which was to developed a natural language database interface [2]. It started to be obvious that meaning needs to be defined contextually and also the subjective aspects is relevant. The word 'red' has different interpretations in different contexts such as "red shirt", "red skin", or "red wine". Subjectivity is particularly notable when abstract complex concepts such as 'computation', 'democracy', or 'sustainability' are considered.

In human communication, it is the occasional clear failure that allows us to see that understanding language is often difficult. In making the connection between a word and its typical and appropriate use, we humans rely on a long learning process. The process is made possible and guided by our genetic make-up, but its success essentially requires extensive immersion to a culture and contexts of using words and expressions. To the extent that these contexts are shared among individual language speakers, we are then able to understand each other. When our learning contexts differ, however, differences in understanding the concepts themselves arise and subsequent communication failures begin to take place. Two main failure types can be detected. The first type is false agreement, where on the surface it looks as if we agree, but in fact our conceptual difference hides the underlying difference in opinions or world views. The second type of problem caused by undiscovered meaning differences is false disagreement. If we are raised (linguistically speaking) in different sub-cultures, we might come to share ideas and views, but might have learned to use different expressions to describe them.

It is commonplace in linguistics to define semantics as dealing with prototypical meanings whereas pragmatics would be associated with meanings in context. For our purposes, this distinction is not relevant since interpretation of natural language expressions always takes place in some context, usually even within multiple levels of context including both linguistic and extra-linguistic ones. In the contrary case, that is, when an ambiguous word such as "break" appears alone without any specific context one can only try to guess which of its multiple meanings could be in question. If there is even a short contextual cue — "break the law", or "have a break", or "how to break in billiards" – it is usually possible to arrive at a more accurate interpretation. Also the extralinguistic context of an expression usually helps in disambiguation.

Becoming conscious of individual differences as a way of increasing understanding

For the most part, people do not seem to be aware of the subjectivity of their perceptions, concepts, or world views. Furthermore, one might claim that we are more typically conscious of differences in opinions, whereas differences in perception or in conceptual level are less well understood. It is even possible that to be able to function efficiently it is best to mostly assume that my tools of communication are shared by people around me. However, there are situations where this assumption breaks to a degree that merits further attention. An example is the case when speakers of the same language from several disciplines, interest groups, or several otherwise closely knit cultural contexts come together to deliberate on some shared issues.

The background assumption of the GICA method innovation is the recognition that although different people may use the same word for some phenomenon, this does not necessarily mean that the conceptualization underlying this word usage is the same; in fact, the sameness at the level of names may hide significant differences at the level of concepts. Furthermore, there may be differences at many levels: experiences, values, understanding of the causal relationships, opinions and regarding the meanings of words. The differences in meanings of words are the most deceptive, because to discuss any of the other differences, a shared vocabulary which is understood in roughly the same way, is necessary. Often a difference in the meanings of used words remains unrecognized for a long time; it may, for instance, be misconstrued as a difference in opinions. Alternatively, a difference in opinions, or regarding a decision that the group makes, may be masked and remain unrecognized, because the same words are used seemingly in accord, but in fact in different meanings by different people. When these differences are not recognized during communication, it often leads to discord and unhappiness about the end result. As a result, the joint process may be considered to have failed in one or even all of its objectives.

Making differences in understanding visible

Our aim with the Grounded Intersubjective Concept Analysis (GICA) method is to devise a way in which differences in conceptualization such as described above can be made visible and integrated into complex communication and decision making processes. An attempt to describe the meaning of one word by relying on other words often fails, because the descriptive words themselves are understood differently across the domains. In fact, a domain may have a large number of words that have their specialized meanings. The more specific aims of this paper are to define the problem domain, to explain the processes of concept formation from a cognitive point of view based on our modeling standpoint, and to propose a methodology that can be used for making differences in conceptual models visible in a way that forms a basis for mutual understanding when different heterogeneous groups interact. Contexts of application are, for instance, public planning processes, environmental problem solving, interdisciplinary research projects, product development processes, and mergers of organizations.

11.5 The GICA method

In the following, we present an overview of the GICA method based on conducting a conceptual survey among participants (see [1] for more details. A version in which text mining can be used to extract subjective information has been introduced recently [3].

The GICA method includes three main stages:

A Preparation,

B Focus session(s), and

C Knowledge to action activities.

These steps can be repeated iteratively. The focus sessions are supported with computational tools that enable the analysis and visualization of similarities and differences in the underlying conceptual systems.

Subjectivity cube

In the GICA method, the idea of considering some items or objects such as words in their contexts is taken a step further. As we have in the introductory section of this paper aimed to carefully show, subjectivity is an inherent aspect of interpretation. In order to capture the aspect of subjectiveness, we add a third dimension to the analysis. Namely, we extend the set of observations, *objects* $\times$ *contexts*, into *objects* $\times$ *contexts* $\times$ *subjects*, i.e. we additionally consider what is the contribution of each subject in the context analysis.

Adopting the notation and terminology for *tensors* (multiway arrays), the *order of a tensor* is the number of the array dimensions, also known as ways or modes. As GICA dataset is observed under varied conditions of three factors, these form the ways of the order-three tensor $\mathcal{X} \in \mathcal{R}^{O \times C \times S}$, where O, C, S are the number of values (levels) in ranges $\{o_1, o_2, \ldots, o_O\}$, $\{c_1, c_2, \ldots, c_C\}$ and $\{s_1, s_2, \ldots, s_S\}$ of the categorical factor variables object $\mathbf{o}$, context $\mathbf{c}$ and subject $\mathbf{s}$ respectively. An element of the tensor, $x_{ijk} \in \mathcal{R}$, is the individual observation under certain values (o_i, c_j, s_k) taken by the factors. $\mathcal{R}$ is the range of the observed variable.[3]

11.5.1 Obtaining subjectivity data

A central question in GICA is how to obtain the data on subjectivity for expanding an object-context matrix into the tensor that accounts additionally for subjectivity. The basic idea is that for each element in the object-context matrix one needs several subjective evaluations. Specifically, *the GICA data collection* measures for each subject s_k the relevance x_{ijk} of an object o_i in a context c_j, or, more generally, the association x_{ijk} between object and context.

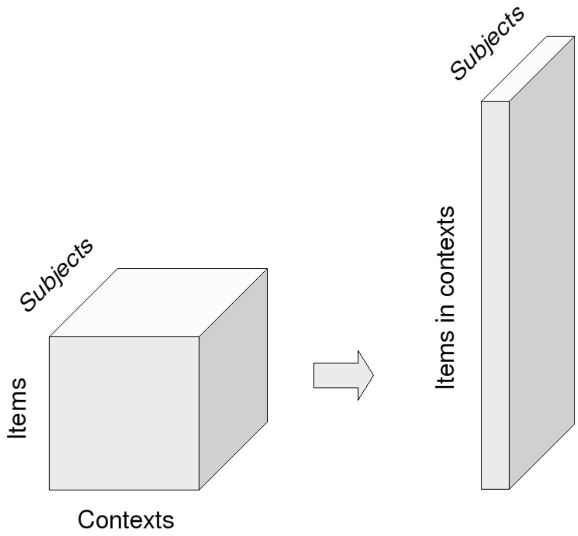

Figure 11.7: The $O \times C \times S$-element subjectivity cube flattened into a matrix in which each column corresponds to a subject and each row to a unique combination of an item and a context. The number of rows in this matrix is $O \times C$ and the number of columns is S. A transpose of this matrix gives rise to a map of persons that is the way-3 matricization of a GICA data tensor.

Conceptual survey of subjectivity

An essential step in the method is to collect a) a number of objects for which epistemological subjectivity is likely to take place, as well as b) a number of relevant contexts towards which the previously collected objects can be reflected. The context items can be short textual descriptions, longer stories, or even multimodal items such as physical objects, images or videos. The underlying idea is that between the objects and the contexts there is some kind of potential link of a varying degree. It is important to choose the contexts in such a manner that they are as clear and unambiguous as possible. The differences in the interpretations of the objects is best revealed if the "reflection surface" of the contexts is as shared as possible among the participants. Therefore, the contexts can include richer descriptions and even multimodal grounding.

The participants are then asked to fill in a data matrix which typically consists of the objects as rows and the contexts as columns. Each individual's task is to determine how strongly an object is associated with a context. A graded scale can be considered beneficial.

The data collected is analyzed using some suitable data analysis method. The essential aspect is to be able to present the rich data in a compact and understandable manner so that the conceptual differences are highlighted.

Text mining of subjectivity

Conducting a conceptual survey requires considerable amount of resources and therefore alternative means for obtaining subjectivity data are useful. As an alternative approach, text mining can be used in this task. The basic idea is to analyze a number of documents stemming from different persons and to compare the use of a set of words or phrases by them. The comparison is based on analyzing the contexts in which each person has used each word. The more similar the contextual patterns between two persons for a word, the closer the conceptions are considered to be. The accuracy of the result is, of course, dependent on how much relevant data is available.

The method is illustrated by analyzing the State of the Union addresses by US presidents. These speeches have been given since president George Washington in 1790. For the detailed analysis, we selected all speeches between 1980 and 2011 given by Jimmy Carter, Ronald Reagan, George Bush, Bill Clinton, George W. Bush and Barack Obama. In this text mining case, populating the matrix takes place by calculating the frequencies on how often a subject uses an object word in the context of a context word. A specific feature in this study was that each president has given the State of the Union Address several times. The basic approach would be to merge all the talks by a particular president together. However, a further extension can be used, i.e., each year can be considered separately so that each president is "split" into as many subjects as the number of talks he has given (e.g. Reagan_{1984}, Reagan_{1985}, ...). This is a sensible option because it provides a chance to analyze the development of the conceptions over time. In our case, the vicinity was defined as 30 words preceding the object word. This contextual window cannot be the whole document because all the objects in a speech would obtain a similar status. On the other hand, a too short window would emphasize the syntactic role of the words. Fig. 11.8 shows a detailed view on the health area of the GICA map. Two specific conclusions can be made. First, a general tendency is that the handling of the health theme forms two clusters, the democrats of the left and the republicans on the right. However, the second conclusion is that in Barack Obama's speeches in 2010 and 2011, he has used the term in such a way that resembles the republican usage.

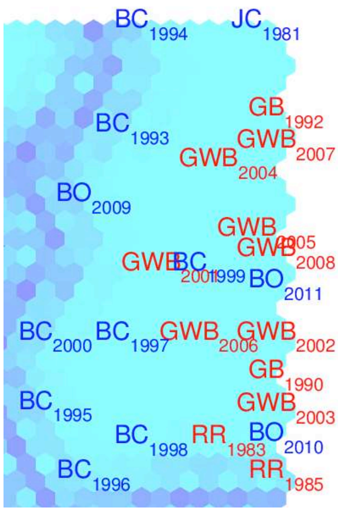

Figure 11.8: A zoomed view into the health area of the GICA map of the State of the Union Addresses.

More detailed information on the GICA method and examples of its use is available in [1, 3]. Future plans include developing the analysis, e.g., by applying different tensor analysis methods.

References

[1] T. Honkela, N. Janasik, K. Lagus, T. Lindh-Knuutila, M. Pantzar, and J. Raitio. GICA: Grounded intersubjective concept analysis – a method for enhancing mutual understanding and participation. TECHREP TKK-ICS-R41, AALTO-ICS, ESPOO, Dec. 2010.

[2] Harri Jäppinen, Timo Honkela, Heikki Hyötyniemi, and Aarno Lehtola. A multilevel natural language processing model. In *Nordic Journal of Linguistics*, 11:69–82, 1988.

[3] Timo Honkela, Juha Raitio, Krista Lagus, Ilari T. Nieminen, Nina Honkela, and Mika Pantzar. Subjects on Objects in Contexts: Using GICA Method to Quantify Epistemological Subjectivity. In *Proceedings of IJCNN 2012, International Joint Conference on Neural Networks*, 2012.