



European Research Council

Established by the European Commission

From Idea to Project: Experiences towards an ERC Advanced Grant

Marie-Francine Moens

KU Leuven

ERC Advanced grants info session

Brussels – April 26, 2019

KU LEUVEN

My background

- 1992: MSc in Computer Science, KU Leuven
- 1999: PhD in Computer Science, KU Leuven
- Since 2015: Full Professor in the Department of Computer Science, KU Leuven
- Main field of expertise: Natural Language Processing
- Research fits ERC PE6: Computer Science and Informatics

2017 ERC Advanced Grant CALCULUS

Commonsense and Anticipation enriched Learning of
Continuous representations supporting **Language
UnderStanding**

H2020-ERC-2017-ADG 788506

What was the motivation to apply for an ERC Advanced Grant?

1. To do fundamental research on a topic I dreamed of for a long time
2. To have more time to focus on high risk research
3. To have time to write a book
4. To obtain recognition and impact of ideas
5. Not to have to spend time to constantly apply for research grants

What counts in the application ?

- CV and track record
- Project

My CV

- Interdisciplinary profile
- Strong track record in attracting funding (e.g., coordinator of FET Open)
- My research team (at the time of application): 1 postdoc and 9 PhD researchers
- My h-index at time of application: 37 (but growing)

Convince your panel that you are at the forefront of your research field (in Europe)
Convince your panel of your key strengths and accomplishments: this might be different for everyone

Track record: some suggestions

Research and innovative thinking: explain the innovations realized in past projects

Independency and leadership: stress own realizations and management capabilities

Leadership in the training of young scientists: what important positions have your PhDs now?

International visibility: e.g., keynotes, tutorials, citations, chair roles in conferences



Track record: some suggestions

	Moons	Part B1	CALDAUS
Publications in top-tier venues	<p>Language 26 (5), 584-605. (41 citations) (Impact factor: 1.903)</p> <p>Gomes, J.C. & Moens, M.-F. (2012). PCA document reconstruction for email classification. <i>Computational Statistics and Data Analysis</i> 56 (3), 741-751. (Impact factor: 1.650) (35 citations)</p> <p>Heyman, S., Vulić, I. & Moens, M.-F. (2016). C-BLDA: extracting cross-lingual topics from non-parallel texts by distinguishing shared from unshared content. <i>Data Mining and Knowledge Discovery</i> 30 (5), 1299-1323. (Impact factor: 2.714)</p> <p>Kalamiridis, G. & Moens, M.-F. (2011). A survey on question answering technology from an information retrieval perspective. <i>Information Sciences</i> 181 (24), 5412-5434. (Impact factor: 4.812) (76 citations)</p> <p>Kordjamehdi, P., Auth, D. & Moens, M.-F. (2015). Structured learning for spatial information extraction from biomedical text: Bacteria biotopes. <i>BMC Bioinformatics</i> 16, 129. (Impact factor: 2.448) (10 citations)</p> <p>Kordjamehdi, P., van Otterlo, M. & Moens, M.-F. (2011). Spatial role labeling: Towards extraction of spatial relations from natural language. <i>ACM Transactions on Speech and Language Processing</i> 8 (3), article 4. (Impact factor: 1.225) (72 citations)</p> <p>Mochales Palau, R. & Moens, M.-F. (2013). Argumentation mining. <i>Artificial Intelligence & Law</i> 19 (3), 1-22. (Impact factor of 2013: 2.36) (124 citations)</p> <p>Pham, P.-F., Moens, M.-F. & Tsykarev, T. (2020). Cross-media alignment of names and faces. <i>IEEE Transactions on Multimedia</i> 22 (1), 14-27. (Impact factor: 3.505) (44 citations)</p> <p>Pham, P.-F., Moens, M.-F. & Tsykarev, T. (2021). Naming persons in news videos with label propagation. <i>IEEE Multimedia</i> 18 (3), 44-55. (Impact factor: 2.849) (44 citations)</p> <p>Vulić, I., De Smet, W., Tang, J. & Moens, M.-F. (2015). Probabilistic topic modeling in multilingual settings: An overview of its methodology and applications. <i>Information Processing & Management</i> 51 (3), 111-147. (Impact factor: 2.391) (29 citations)</p> <p>Vulić, I. & Moens, M.-F. (2016). Bilingual distributed word representations from document-aligned comparable data. <i>Journal of Artificial Intelligence Research</i> 55, 953-994. (Impact factor 2016/2017: 2.334) (31 citations).</p> <p>Zoghbi, S., Vulić, I. & Moens, M.-F. (2016). Latent Dirichlet allocation for linking user-generated content and e-commerce data. <i>Information Sciences</i> (Impact factor: 4.812) (2 citations).</p>		
International collaborations with top-researchers and institutes	<p>I collaborate and have collaborated with numerous European research teams as well as with Chinese groups. My team regularly welcomes foreign guests (academic staff, postdoctoral fellows and PhD students). I encourage my researchers to make research visits at foreign institutes. For instance, three PhD students (Thomas Provost, Susana Zoghbi and Golnoosh Farnadi) visited for several months Tsinghua University in China, and Susana Zoghbi and Golnoosh Farnadi have visited Microsoft in Cambridge, UK and Seattle, USA respectively. Current postdoctoral researcher Susana Zoghbi has recently completed an internship at NASA, USA, and PhD researcher Guillem Collet is completing an internship at ETH Zurich.</p>		
Entrepreneurship and collaboration with international companies	<p>Entrepreneurship and collaboration with companies</p> <p>In my research projects, I have collaborated with important companies such as Symantec, Tiscali, Thales Security Systems, Telefonica Investigacion y Desarrollo, Atos Origin and Space Applications Services among others resulting in patent applications. I find it important to work on a mixture of fundamental and more application oriented projects.</p>		
Teaching and scientific services: shows expertise and recognition in the field	<p>Teaching</p> <p>My favorite courses are Natural Language Processing and Text Based Information Retrieval which attract each year a growing number of students. I am a regular speaker at summer schools and during tutorials of international conferences.</p> <p>Scientific services</p> <p>I was chair of the European Chapter of the Association for Computational Linguistics and member of the executive council of the Association for Computational Linguistics in 2011 and 2012. As program chair, area chair and technical co-chair of top conferences and member of the editorial board of journals in the domains of Natural Language Processing, Information Retrieval and Machine Learning, I help actively in the distribution of research.</p> <p>Motivation for the application for the ERC Advanced grant</p> <p>The grant would give me the opportunity to work on the research topic of automated language understanding in which I have utmost interest since my master studies at UCLA.</p>		

The project

- **The most important thing**
- **Key idea of the proposal must:**
 - Be creative and innovative
 - Be high risk / high gain
 - Be relevant
 - Have a potential of high impact
- **You must be the ideal person for the job**

Key idea of the CALCULUS project

Current **neuroscience studies** show that humans perform language understanding tasks instantly by relying on their capability to imagine situations.

The **goal of CALCULUS** is to study **new paradigms for machines that learn to understand language**, for fast parsing of language and efficient use of memory, and for inference in imagined and real physical spaces that are inspired by human intelligence.



What makes the project high risk/high gain?

- CALCULUS: **paradigm shift:**
 - Since the early days of AI (from 1956 onwards): language understanding = translation to discrete symbolic representations and reasoning with these \Rightarrow poor representation of reality
- CALCULUS:
 - Translation of language to continuous representations that integrate (visual) context
 - Reasoning in 3D/4D imagined or real-world spaces, or reasoning with the continuous representations that generate these spaces

What makes the project high risk/high gain?

At the time of submission of the proposal in 2017:

- Insights into the real bottlenecks of language understanding from a previous FET-open project

Shows real good understanding of the problems

- A few publications of our lab in highly ranked conferences that support the novel paradigm

Shows that objectives are achievable

CALCULUS = novel machine learning paradigms inspired by results of neuroscience research with **impact** on:

Impact



<https://www.quotemaster.org/Robots>

Machine reading of
text



Human-robot natural
language dialogue



Indexing, storage and
retrieval of continuous
representations of
content



<https://alsadotorg.wordpress.com/2016/06/02/bringing-brain-computer-interface-home/>

Brain-computer
interfaces

Natural language
processing

Natural language
processing and robotics

Information retrieval

Human-computer interaction

Takeaway

- Pose innovative, high-impact research hypotheses and objectives
- Propose a novel solution to a long-standing problem

And provide preliminary evidence (obtained by you or other researchers) to support these: contributes to feasibility

- Show that now is the right time to investigate the proposed solution
- Convince the panel that you are the ideal person for the job:
 - Expertise
- Clarify the impact also beyond your field

How was the writing done?

1. First quick draft with ideas (submission in 2015)
2. Refinement with many technical details and literature studies including the current neuroscience studies (submission in 2016)
3. Final version: only few content changes were made, but technical details were better framed, explained and motivated; **our own preliminary results added** (submission in 2017)

Methodology section: novel approaches with good motivation and source of inspiration

<p>2.3. Inference and learning of new knowledge</p> <p>Source of inspiration</p> <p>Humans (Vernon 2004 p. 146 ff., Downing 2005 p. 231 ff.) learn with supervision of a teacher, through reinforcement by trial and error, and without any supervision just by exploration. Language comprehension by humans for a large part relies on the prediction of content. However, when content is described in a text that is incompatible with the imagined content, a prediction error is produced and additional processing is needed (shown by increased fMRI activation in Broca's region of the brain) (Kristensen & Wallentin 2015), and leading to <i>perceptual experiences</i> (Moulton & Kosslyn 2009).</p> <p>2.2. Real time memory based retrieval and semi-supervised parsing</p> <p>Source of inspiration</p> <p>When humans anticipate content while interpreting language, the brain generates the right representations in working memory and makes the inferences that fit the immediate task or context (Lambon Ralph et al. 2017). When interpreting language, some brain scientists argue that local chunks of linguistic structure are incrementally processed (Kristensen & Wallentin 2015).</p> <p>Methodology</p> <p>Although there is project on two anticipatory representations, the latter hierarchically organizes current text input (see above works and representations investigate this representation).</p> <p>Search of most plausible representations</p> <p>First, as traditionally done, we search structures including words, but phrases and information should be next-word indices. This relevant anticipatory representation is very unique and leads to a second option is to train representation to perform. This entails using the</p>	<p>5. Methodology</p> <p>5.1. Learning anticipatory representations from language and visual data</p> <p>Source of inspiration</p> <p>Cognitive and neuroscience studies have demonstrated that the human brain uses anticipation to perform tasks very efficiently (Vernon 2004 p. 2). Based on lifelong verbal and perceptual experiences a human anticipates what events might occur in his or her environment. A similar approach is witnessed when humans read texts (Kutney & Zacks 2011, Lambon Ralph et al. 2017). A reader's comprehension system continuously makes predictions about what information will be presented next in the text. For humans, perceptual experiences play an important role in remembering the past and imagining possible future experiences (Schacter & Madore 2006), which often might involve some degree of scene reconstruction with key elements (people, objects, settings and their relations). These studies confirm older theories that have modelled comprehension as perceived or imagined events (Johnson-Laird 1982). Multiple forms of mental imagery exist, for instance, object-based (of shapes, colors), spatial (e.g., of locations), auditory and motoric (Moulton & Kosslyn 2009), but it is assumed that they rely on perception among which visual experiences play an important role. Neuroscientists demonstrate the existence of large scale modality independent conceptual representations in the brain, apart from small scale modality dependent and category specific representations (Pradipras et al. 2018). They also acknowledge that the anticipated events are stored as template structures that are independent of the various distinctive contexts that might be encountered, providing a basis for conceptual generalization (Lambon Ralph et al. 2017).</p> <p>Methodology</p> <p>Learning of the representations from multimodal visual and textual data and level of supervision : We believe that learning representations that capture imagined events, their composing parts and their relationships need to be learned from image data paired with language descriptions (Robens 2018). In CALIGULI, we focus on representations that predict semantic frames, that is, actions and their arguments (semantic roles, temporal and spatial relations), and scripts of events or actions. The multimodal data on which the templates will be trained will vary with respect to the degree of support by predefined or preprocessed annotations. In a first instance, we will train on rather structured data in which objects and their relations are already marked (such as in the Visual Genome dataset²). Later we will experiment with weakly structured data with a ground detected bounding boxes (images) and detected entities/actions (text).</p>
--	--

How was the writing done?

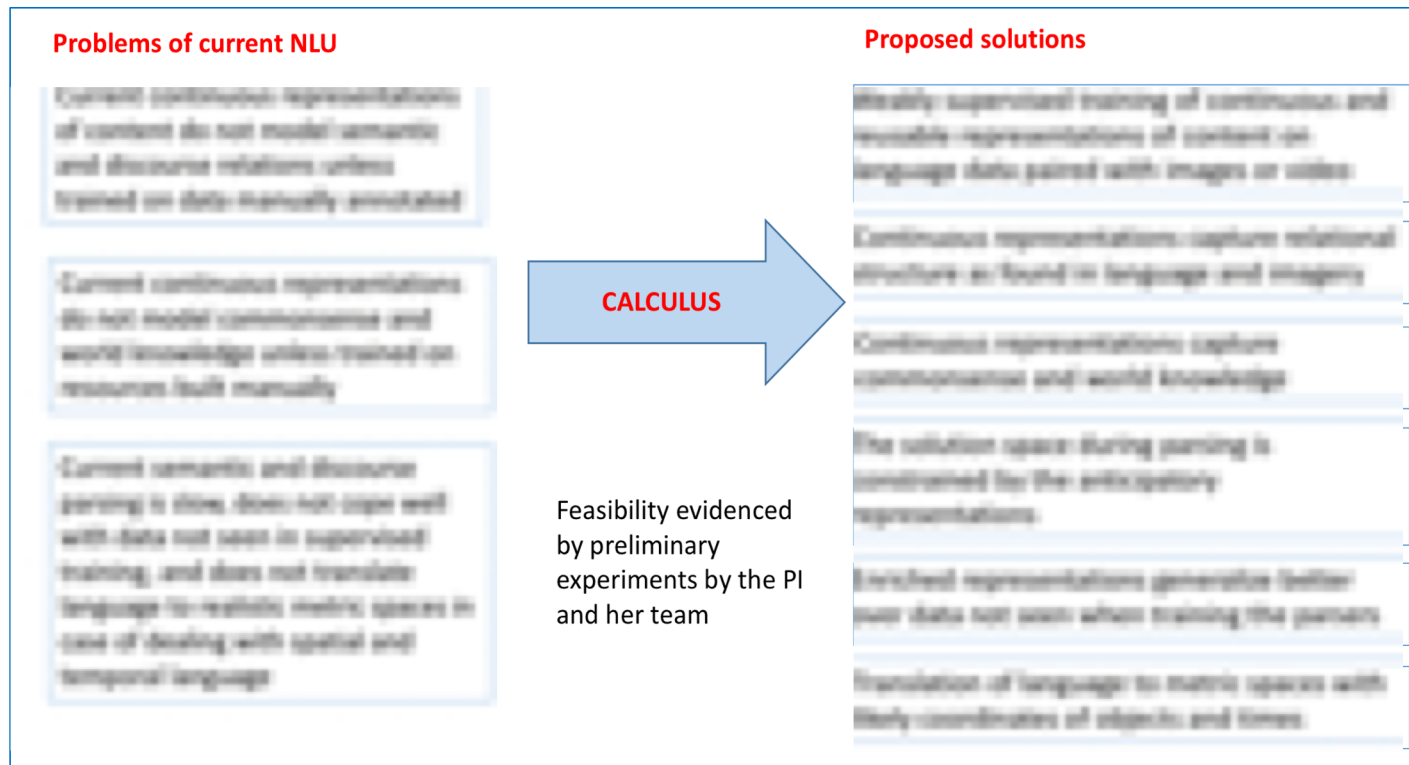


Figure 1. Gaps in current natural language understanding and solutions proposed by CALCULUS.

How was the writing done?

- Make sure that the proposal is written with both the expert and not-so-expert in mind
 - Part B1 is reviewed by panel, Part B2 by external (expert) reviewers
- Make your proposal visually attractive, e.g., nice layout and add some pictures
- Have the proposal read by colleague researchers

Conclusion

- If you have a track record of excellent science
- If you have an innovative idea (that preferably has some interdisciplinary flavor)
- If you can somehow demonstrate that this idea will work

=> Dream, think, dare, write an excellent proposal and do not give up

Thank you and success !



European Research Council

Established by the European Commission