

# Dr. Stevan Rudinac

University of Amsterdam  
Faculty of Economics and Business  
Amsterdam Business School

28 November 2025

T: +31 (0)20 525 5254

E: s.rudinac@uva.nl

W: www.stevanrudinac.com

## Education

• <b>Delft University of Technology</b>	Delft, NL
• <i>PhD Computer Science, field - Multimedia Information Retrieval</i>	2008 - 2013
• <b>Faculty of Electrical Engineering, University of Belgrade</b>	Belgrade, SRB
• <i>Graduate Engineer of Electrical Engineering</i>	2001 - 2006
• <b>Military Technical Academy</b>	Belgrade, SRB
• <i>Air Technical Service - Electronic Engineering</i>	2000 - 2001
• <b>Military Grammar School</b>	Belgrade, SRB
• <i>Gymnasium - secondary school</i>	1996 - 2000

## Work Experience

• <b>University of Amsterdam, Amsterdam Business School</b>	Amsterdam, NL
• <i>Associate Professor</i>	2019 - present
• <b>University of Amsterdam, Informatics Institute</b>	Amsterdam, NL
• <i>Guest Researcher</i>	2017 - present
• <b>University of Amsterdam, Multimedia Information Retrieval for Business Lab</b>	Amsterdam, NL
• <i>Scientific Director</i>	2022 - present
• <b>University of Amsterdam, Amsterdam Business School</b>	Amsterdam, NL
• <i>Assistant Professor</i>	2017 - 2019
• <b>University of Amsterdam, Informatics Institute</b>	Amsterdam, NL
• <i>Postdoctoral Researcher</i>	2013 - 2017
• <b>Netherlands Forensic Institute</b>	Den Haag, NL
• <i>Forensic Scientist</i>	2012 - 2013
• <b>Delft University of Technology</b>	Delft, NL
• <i>PhD Researcher</i>	2008 - 2012
• <b>Eindhoven University of Technology</b>	Eindhoven, NL
• <i>PhD Researcher</i>	2007 - 2008
• <b>ICEF, Faculty of Electrical Engineering, University of Belgrade</b>	Belgrade, SRB
• <i>Research Assistant</i>	2006 - 2007

## Recent Teaching Experience

• <b>Coordinator: MSc DS&amp;BAN course Applied AI Research Seminar</b>	UvA
• <i>Topic: deep learning for computer vision, multimedia and information retrieval</i>	2022 - present
• <b>Lecturer and coordinator: MBA Big Data course Deep Learning</b>	UvA
• <i>Topic: Introduction to deep learning</i>	2021 - present
• <b>Lecturer and coordinator: MBA Big Data course Machine Learning</b>	UvA
• <i>Topic: Introduction to machine learning in Python</i>	2018 - present
• <b>Lecturer and coordinator: Minor DS course Machine Learning</b>	UvA
• <i>Topic: Introduction to machine learning in Python</i>	2017 - present
• <b>Lecturer and coordinator: Minor DS course Data Structures and Algorithms</b>	UvA
• <i>Topic: Common algorithms, algorithmic paradigms, and data structures</i>	2017 - present

• <b>Lecturer: BSc Information Studies course Information Visualisation (Dutch)</b>	UvA
• <i>Topic: Principles and techniques of information visualization</i>	2014 - 2017
• <b>Lecturer and Guest Lecturer: MSc IS course Fundamentals of Data Science</b>	UvA
• <i>Topic: Hands on experience in data gathering, analysis and visualisation</i>	2016, 2018, 2019
• <b>Lecturer: Lean Six Sigma Master Black Belt Masterclass on Machine Learning</b>	UvA
• <i>Topic: "A tour through the machine learning zoo"</i>	2018, 2019
• <b>Guest lecturer: BSc IS course Onderzoek van Interactieve Systemen (Dutch)</b>	UvA
• <i>Topic: Semantic multimedia analysis for researching violent online political extremism</i>	2018 - 2022
• <b>Guest lecturer: MSc Information Studies course Intelligent Interactive Systems</b>	UvA
• <i>Topic: Multimedia analytics applications in urban computing</i>	2014 - 2020

## Supervision of PhD and MSc Students

• <b>AI for sustainable investment</b>	ABS, UvA
• <i>Marleen de Jonge: Supervisor (with A. Mićković &amp; S.J.G. van Wijnbergen)</i>	2024 - present
• <b>Human-aligned video-AI for responsible business</b>	DSC HAVA-Lab: ABS & Ivi, UvA
• <i>Gowreesh Mago: Co-supervisor (with P.S.M. Mettes &amp; C.G.M. Snoek)</i>	2024 - present
• <b>Network statistics and machine learning</b>	ABS, UvA
• <i>Yunming Hui: Co-supervisor (with I.M. Zwetsloot and S. Trimborn)</i>	2023 - present
• <b>Multimodal conversational search for business</b>	ABS, UvA
• <i>Hongyi Zhu: Supervisor (with E. Kanoulas, J. Demmers and D.N. den Hartog)</i>	2022 - present
• <b>Multimodal hypergraph learning for cultural industries and science</b>	ABS, UvA
• <i>Shuai Wang: Co-supervisor (with M. Kackovic, N. Wijnberg and M. Worring)</i>	2022 - present
• <b>Tipping points in categorization and canonization in the art market</b>	ABS, UvA
• <i>Thanos Eftymiou: Co-supervisor (with M. Kackovic, N. Wijnberg and M. Worring)</i>	2020 - present
• <b>Measuring the impact of high performance computing in healthcare</b>	ABS, UvA
• <i>Jieyi Li: Supervisor (with A. Osseyrany and A. Berg):</i>	completed 2025
• <b>AI and Marketing: Analyzing abstract dimensions in online social multimedia</b>	ABS, UvA
• <i>Ujjwal Sharma: Supervisor (with W.M. van Dolen, M. Worring and J. Demmers)</i>	2018 - present
• <b>Multimodal machine learning for addressing grand societal challenges</b>	ABS & Ivi, UvA
• <i>Maarten Sukkel: Supervisor (with M. Worring)</i>	completed 2024
• <b>Multimodal deep learning on hypergraphs</b>	Ivi, UvA
• <i>Devanshu Arya: Supervisor (with M. Worring)</i>	completed 2022
• <b>The Machine in multimedia analytics</b>	Ivi, UvA
• <i>Jan Zahálka: Co-supervisor (with M. Worring)</i>	completed 2017
• <b>Supervisor of ~ 100 MSc students</b>	UvA
• <i>Topics: multimedia, computer vision, information retrieval, machine learning</i>	2013 - present

## Outreach: Recent Public Talks, Seminars and Interviews

• <b>Keynote: From Content-Based Image and Video Retrieval With Relevance Feedback to Multimodal Large Language Models and Beyond: A Journey Through Modalities and Semantic Levels</b>	Dublin, IE
• <i>ACM MM 2025 Workshop on Multimedia Analytics with Multimodal Large Language Models</i>	2025
• <b>Penelist: Future directions in multimedia analytics with MLLMs</b>	Dublin, IE
• <i>ACM MM 2025 Workshop on Multimedia Analytics with Multimodal Large Language Models</i>	2025
• <b>Panelist: Multimodal Content Generation and Moderation</b>	Nara, JP
• <i>MMM'25 Special Session on Multimodal Large Language Models and Applications</i>	2025
• <b>Panelist: Perspectives on how AI can foster a more sustainable future</b>	Amsterdam, NL
• <i>Amsterdam AI Highlights Event</i>	2024

<b>Talk: Hoe de inzet van multimedia computing kan helpen de leefbaarheid van de stad te verbeteren</b>	Amsterdam, NL
<i>Urban Impact: samenwerking UvA en Gemeente Amsterdam</i>	2024
<b>Co-chair: 2021 Amsterdam Data Science Thesis Awards</b>	Amsterdam, NL
<i>Promoting excellence in Data Science and AI in all Amsterdam-based knowledge institutes.</i>	2021
<b>Talk: Multimodal Representation Learning</b>	Amsterdam, NL
<i>Amsterdam Data Science (ADS) &amp; Prosus Webinar on Universal Representation</i>	2021
<b>Moderator of the Webinar “A-Z in Data Science: Computer Vision”</b>	Amsterdam, NL
<i>Amsterdam Data Science (ADS) &amp; SABA Webinar Series</i>	2020
<b>Keynote: How multimedia analytics can help solve complex business problems.</b>	Kiev, UA
<i>9th Data Science UA Conference (online)</i>	2020
<b>Panellist in a Panel Discussion: Multimedia, not Single Media...</b>	Dublin, IE
<i>International Conference on Content-Based Multimedia Indexing, CBMI 2019</i>	2019
<b>Keynote: Multimedia Analytics for Business</b>	Kiev, UA
<i>Data Science UA - 7th Conference on Machine Learning, Artificial Intelligence and Data Science</i>	2019
<b>Interview: In gesprek met de stad. Booklet “De Universiteit en Amsterdam”</b>	Amsterdam, NL
<i>Article published on the occasion of the opening of the academic year</i>	2018
<b>Talk: Multimedia Analytics for Urban Computing</b>	Amsterdam, NL
<i>ADS Meetup “Coffee &amp; Data: Urban Analytics”</i>	2018
<b>Talk: Data Science in Service of City Liveability</b>	Amsterdam, NL
<i>ADS Meetup “Smart Cities &amp; Smart Health in Amsterdam &amp; Beijing”</i>	2018
<b>Keynote: “Good marketing” use-case applications through open data and social multimedia.</b>	Amsterdam, NL
<i>Innovators Tribe: AI &amp; Multimodal Data, Startup Village, Science Park</i>	2018
<b>Member “Klankbordgroep DrukteRadar van de Gemeente Amsterdam”</b>	Amsterdam, NL
<i>An expert group formed to help city administrators reduce the crowd in Amsterdam</i>	2018 - present
<b>Member “Klankbordgroep voorspelmodel balans in de stad van de Gemeente Amsterdam,”</b>	Amsterdam, NL
<i>An expert group advising city administrators on predictive analysis tools</i>	2017 - 2018
<b>Jury member of the “Entrepreneurship in Data Science” programme</b>	Amsterdam, NL
<i>ACE &amp; ADS programme for MSc &amp; PhD Students, Startup Village, Science Park</i>	2018
<b>User-Generated Content in Social Media</b>	Saarbrücken, DE
<i>Invited to participate in the prestigious Schloss Dagstuhl Seminar</i>	2017
<b>Keynote: Understanding Violent Online Political Extremism</b>	Arnhem, NL
<b>Using Multimedia Analytics</b>	2016
<i>Themadagen digitale opsporing, a major national digital forensics event</i>	2016
<b>Human Reasoning Meets Machine Learning: Hands-on Workshop</b>	Dublin, IE
<b>Using Mixed-Media Analytics Tools on Violent Online Extremism Data</b>	2016
<i>VOX-Pol Mid-Project Conference</i>	
<b>Semantic Multimedia Analysis for Researching Violent Online Political Extremism Content</b>	Den Haag, NL
<i>Workshop at the Security Research &amp; Innovation Event</i>	2016
<b>Talk: Multimedia Analytics Tools for Research in Violent Online Political Extremism</b>	Amsterdam, NL
<i>ADS Meetup “Coffee &amp; Data: Forensic Intelligence meets Data Science”</i>	2016
<b>Talk: Using Social Multimedia for Urban Computing</b>	Amsterdam, NL
<i>Amsterdam Data Science Meetup “Coffee &amp; Data #2”</i>	2015

## Research Grants

- Co-Recipient Research Priority Area (RPA)**
- Building Interdisciplinary Team Science for Future Proof Higher Education*

UvA  
2025 - present

- Co-applicants: Ingmar Visser (FMG/Psy), Niels Smits (FMG/POW), Peter van Baalen (FEB) Tom Broens (FdG/AUMC); Jelke Bloem (FGw); Erwin van Vliet (FNWI); Monique Volman (FMG/POW); Suzanne Geerlings (FdG/AUMC); Stevan Rudinac (FEB); Brigitte ten Brink (FMG/Psy); Mieke Lopes Cardozo (FMG/SW; GPIO); Edwin van Hooft (FMG/Psy); Thijs Bol (FMG/Soc); Frank Cornelissen (FMG/POW); Natalie Pareja Roblin (FMG/POW); Joost Jansen in de Wal (FMG/POW); Brenda Jansen (FMG/Psy); Ana Oprescu (FNWI); Josje Verhagen (FGw); Sible Andringa (FGw); Inez Zwetsloot (FEB)
- Funding: 1.500.000 euro

- **Co-Recipient/Co-Principal Investigator UvA DSC Lighthouse Grant** UvA  
  - *Human-Aligned Video AI Laboratory (DSC HAVA-Lab)* 2023 - present
    - Co-applicants: Prof.dr. C. Snoek (Lead Principal Investigator), dr. P. Mettes (Lab Manager), dr. E. Berkhout, Prof.dr. T. Blanke, dr. I. Groen, dr. mr. H. Janssen, Prof.dr. M. Lindegaard, and Prof.dr. M. Schijven
    - Funding: 7 PhD Students (1 PhD student supervised by me and dr. P. Mettes)
- **Co-recipient of a UvA/a.s.r./TKI ICT grant** UvA  
  - *Topic: Research Centre for Sustainable Investments & Insurance (RCSII)* 2023 - present
    - Funding: Total funding of 2.767 million euro, including 308.000 from TKI ICT PPS; 1 PhD and 1 Postdoc under my supervision.
- **Co-recipient of a UvA DSC Interdisciplinary Ph.D. Program grant** UvA  
  - *Topic: Discovering the secret ingredients of successful innovations* 2021 - present
    - Co-applicants: dr. M. Kackovic, Prof.dr. N. Wijnberg and Prof.dr. M. Worring
    - Funding: PhD Student
- **Recipient of a Google TensorFlow grant for support in teaching ML/DL courses** UvA  
  - *Topic: A Tour Through the Deep Learning Zoo* 2021 - present
    - Funding: \$8.000
- **Co-recipient of a UvA/FEB research grant** UvA  
  - *Topic: Tipping points in categorization and canonization in the art market* 2019 - present
    - Partners: Sotheby's (New York), Rijksmuseum, Stedelijk Museum Amsterdam, and the Netherlands Institute for Sound and Vision
    - Co-applicants: dr. M. Kackovic, Prof.dr. N. Wijnberg and Prof.dr. M. Worring
    - Funding: PhD Student
- **Recipient of an NVIDIA Data Science GPU Grant** UvA  
  - *Topic: Multimedia Analytics for Urban Computing* 2019 - 2022
    - Funding: Refurbished NVIDIA Titan V GPU
- **Co-recipient of Amsterdam Business School research grant** UvA  
  - *Topic: Artificial Intelligence and Marketing* 2018 - present
    - Partners: Starbucks
    - Co-applicants: Prof.dr. Willemijn van Dolen
    - Funding: PhD Student
- **Co-recipient of a research grant from the City of Amsterdam** UvA  
  - *Topic: Improving Response to Citizen Reports Using Artificial Intelligence* 2018 - 2024
    - Partners: Chief Technology Office and DataLab, City of Amsterdam
    - Co-applicants: Prof.dr. Marcel Worring
    - Funding: PhD Student, two days per week
- **Co-recipient of a research grant within EU Interreg SCORE project** UvA  
  - *Topic: Smart Cities + Open Data Re-use (SCORE)* 2017 - 2022
    - Partners: Nine cities from the EU North Sea Region
    - Co-applicants: Prof.dr. Marcel Worring
    - Funding: PhD Student

- **Co-recipient of Amsterdam Data Science research grant** UvA  
2016 - 2017
  - *Topic: Measuring and Improving City Livability*
    - Partners: Amsterdam University of Applied Sciences
    - Funding: Two Research Assistants
  
- **Co-recipient of Amsterdam Data Science research grant** UvA  
2015 - 2016
  - *Topic: CROWDEO - Human-Machine Video Summarization*
    - Partners: VU Amsterdam
    - Funding: Research Assistant

## Academic Service

- General Chair of the 16th ACM International Conference on Multimedia Retrieval, ACM ICMR 2026, Amsterdam, The Netherlands
- Technical Program Chair of the 32nd International Conference on MultiMedia Modeling, MMM 2026, Prague, Czech Republic
- Technical Program Chair of the 33rd ACM International Conference on Multimedia, ACM Multimedia 2025, Dublin, Ireland
- Workshop Chair of the 6th ACM International Conference on AI in Finance, ICAIF 2025, Singapore
- General Chair of the 30th Int. Conference on MultiMedia Modeling, MMM 2024, Amsterdam, Netherlands
- Technical Program Chair of the 29th Int. Conference on MultiMedia Modeling, MMM 2023, Bergen, Norway
- Invited speakers and tutorial chair of the 19th International Conference on Content-based Multimedia Indexing, CBMI 2022, Graz Austria
- Panel Chair of the 29th ACM International Conference on Multimedia, ACM Multimedia 2021, Chengdu, China, hybrid (onsite and virtual) event
- Co-organizer of the 1st International Workshop on Multimedia Computing for Urban Data. In Conjunction with ACM Multimedia 2021, Chengdu, China, hybrid (onsite and virtual) event
- Lead Guest Editor of Multimedia Tools and Applications (MTAP) journal Special Issue “Content-Based Multimedia Indexing in the Era of Artificial Intelligence”. Publisher: Springer Nature. May 2021.
- Co-organizer of special session on Multimedia Information Retrieval for Urban Data. In 10th International Conference on Multimedia Retrieval, ACM ICMR 2020, Dublin, Ireland.
- Associate Editor/Area Chair of the 25th Int. Conference on Pattern Recognition, ICPR 2020, Milan, Italy
- Program Chair of the 2019 International Conference on Content-Based Multimedia Indexing, Dublin, Ireland
- Lead Guest Editor of Multimedia Tools and Applications (MTAP) journal upcoming Special Issue “Content-Based Multimedia Indexing in the Era of Artificial Intelligence: Special Issue of CBMI 2019”
- Co-organizer of special session on Multimedia Analytics: Perspectives, Techniques and Applications (MAPTA). In:
  - 24th Int. Conf. on Multimedia Modeling (MMM '18), February 5-7, 2018, Bangkok, Thailand
  - 25th Int. Conf. on Multimedia Modeling (MMM '19), January 8-11, 2019, Thessaloniki, Greece
  - 26th Int. Conf. on Multimedia Modeling (MMM '20), January 5-8, 2020, Daejeon, Korea
  - 27th Int. Conf. on Multimedia Modeling (MMM '21), January 25-27, 2021, Prague, Czech Republic
- Remote expert Evaluator for the European Commission’s Horizon 2020 FET-Open RIA calls, 2018-present
- Member of the VOX-Pol Network of Excellence (NoE), a European Union Framework Programme 7 (FP7)-funded academic research network focused on researching the prevalence, contours, functions, and impacts of Violent Online Political Extremism and responses to it, [www.voxpol.eu](http://www.voxpol.eu), 2015-2018.
- Member of the organizing committee of ACM Multimedia 2016, Amsterdam, The Netherlands
- Co-organiser of CrowdMM 2015 Workshop at ACM Multimedia 2015, Brisbane, Australia
- Organizing committee of S3MR '11, 2nd Summer School on Social Media Retrieval, Antalya, Turkey
- Organizing committee of S3MR '10, 1st Spring School on Social Media Retrieval, Interlaken, Switzerland
- Served as reviewer/programme committee member for a number of international journals and conferences in multimedia and information retrieval fields, such as:
  - JOURNALS: IEEE Transactions on Multimedia (IEEE TMM) | ACM Transactions on Multimedia Computing, Communications, and Applications (ACM TOMM) | IEEE MultiMedia Magazine (IEEE MM) | ACM Transactions on Information Systems (ACM TOIS) | IEEE Transactions on Visualization

and Computer Graphics (IEEE TVCG) | Information Retrieval Journal | Forensic Science International: Digital Investigation | Information Processing & Management

- CONFERENCES: ACM International Conference on Multimedia (ACM MM) | ACM International Conference on Multimedia Retrieval (ACM ICMR) | IEEE International Conference on Multimedia and Expo (IEEE ICME) | European Conference on Information Retrieval (ECIR) | ACM International Conference on Information and Knowledge Management (ACM CIKM).

## Honors and Awards

• <b>Lecturer of the Year at the University of Amsterdam Business School</b>	ABS, UvA
• <i>Machine Learning and Deep Learning Courses in the MBA Big Data &amp; Business Analytics</i>	2022
• <b>UvA Economics and Business Lecturer of the Year</b>	UvA
• <i>Faculty Lecturer of the Year at the UvA Lecturer of the Year 2021 Awards</i>	2021
• <b>Van der Schroeff Award for the best lecturer at UvA Economics and Business</b>	EB, UvA
• <i>Annual award for the faculty lecturer of the year</i>	2020
• <b>Lecturer of the Year at the University of Amsterdam Business School</b>	ABS, UvA
• <i>Machine Learning Course in the MBA Big Data &amp; Business Analytics</i>	2019
• <b>Grand Challenge 1st Prize at ACM Multimedia 2014 Conference</b>	Orlando, FL, USA
• <i>with Jan Zahálka and Marcel Worring</i>	2014
• <b>Outstanding Reviewer Award at the Pacific-Rim Conference on Multimedia</b>	Kuching, MY
• <i>Awarded for producing high-quality and timely reviews</i>	2014
• <b>Graduated as first in class from the Military Grammar School in Belgrade</b>	Belgrade, SRB
• <i>Annual award for the best graduate of the grammar school</i>	2000

## Publications (Since 2010)

- [1] Devanshu Arya, Deepak K. Gupta, Stevan Rudinac, and Marcel Worring. Adaptive neural message passing for inductive learning on hypergraphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(1):19–31, 2025.
- [2] Yunming Hui, Inez Maria Zwetsloot, Simon Trimborn, and Stevan Rudinac. Domain-informed negative sampling strategies for dynamic graph embedding in meme stock-related social networks. In *Proceedings of the ACM on Web Conference 2025*, WWW '25, page 518–529, New York, NY, USA, 2025. Association for Computing Machinery.
- [3] Jia-Hong Huang, Yixian Shen, Hongyi Zhu, Stevan Rudinac, and Evangelos Kanoulas. Gradient weight-normalized low-rank projection for efficient llm training. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(23):24123–24131, Apr. 2025.
- [4] Athanasios Efthymiou, Monika Kackovic, Stevan Rudinac, Marcel Worring, and Nachoem Wijnberg. Being ranked in a material world: The visual originality of an artwork and its effects on the artist’s canonization. *Organization Studies*, 2025. Accepted for Publication.
- [5] Hongyi Zhu, Jia-Hong Huang, Yixian Shen, Stevan Rudinac, and Evangelos Kanoulas. Interactive image retrieval meets query rewriting with large language and vision language models. *ACM Trans. Multimedia Comput. Commun. Appl.*, 21(10), October 2025.
- [6] Shuai Wang, Ivona Najdenkoska, Hongyi Zhu, Stevan Rudinac, Monika Kackovic, Nachoem Wijnberg, and Marcel Worring. Artrag: Retrieval-augmented generation with structured context for visual art understanding. In *Proceedings of the 33rd ACM International Conference on Multimedia*, MM '25, page 6700–6709, New York, NY, USA, 2025. Association for Computing Machinery.
- [7] Luca Rossetto, Werner Bailer, Duc-Tien Dang-Nguyen, Graham Healy, Björn Pór Jónsson, Onanong Kongmeesub, Hoang-Bao Le, Stevan Rudinac, Klaus Schöffmann, Florian Spiess, Allie Tran, Minh-Triet Tran, Quang-Linh Tran, and Cathal Gurrin. The castle 2024 dataset: Advancing the art of multimodal understanding. In *Proceedings of the 33rd ACM International Conference on Multimedia*, MM '25, page 12629–12635, New York, NY, USA, 2025. Association for Computing Machinery.
- [8] Omar Shahbaz Khan, Ujjwal Sharma, Gonçalo Marcelino, Aaron Duane, Stevan Rudinac, Marcel Worring, and Björn Pór Jónsson. Interactive retrieval system for multi-stream collections: multiview at castle 2025 interactive grand challenge. In *Proceedings of the 33rd ACM International Conference on Multimedia*, MM '25, page 14273–14279, New York, NY, USA, 2025. Association for Computing Machinery.

[9] Ujjwal Sharma, Omar Shahbaz Khan, Stevan Rudinac, and Björn Pór Jónsson. Can relevance feedback, conversational search and foundation models work together for interactive video search and exploration? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3749–3758, June 2025.

[10] Luca Rossetto, George Awad, Werner Bailer, Cathal Gurrin, Björn Pór Jónsson, Jakub Lokoč, Stevan Rudinac, and Klaus Schoeffmann. Overview of the 1st international workshop on interactive video search and exploration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3682–3687, June 2025.

[11] Ujjwal Sharma, Omar Shahbaz Khan, Stevan Rudinac, and Björn Pór Jónsson. Exquisitor at the video browser showdown 2025: Unifying conversational search and user relevance feedback. In *MultiMedia Modeling - 31st International Conference on Multimedia Modeling, MMM 2025, Nara, Japan, January 8-10, 2025, Proceedings, Part V*, volume 15524 of *Lecture Notes in Computer Science*, pages 264–271. Springer, 2025.

[12] Jia-Hong Huang, Hongyi Zhu, Yixian Shen, Stevan Rudinac, and Evangelos Kanoulas. Image2text2image: A novel framework for label-free evaluation of image-to-text generation with text-to-image diffusion models. In *MultiMedia Modeling - 31st International Conference on Multimedia Modeling, MMM 2025, Nara, Japan, January 8-10, 2025, Proceedings, Part IV*, volume 15523 of *Lecture Notes in Computer Science*, pages 413–427. Springer, 2025.

[13] Maarten Sukel, Stevan Rudinac, and Marcel Worring. Multimodal temporal fusion transformers are good product demand forecasters. *IEEE MultiMedia*, 31(2):48–60, 2024.

[14] Ujjwal Sharma, Stevan Rudinac, Joris Demmers, Willemijn van Dolen, and Marcel Worring. From pixels to perceptions: Capturing high-level abstract concepts in visual user-generated content. *International Journal of Information Management Data Insights*, 4(2):100269, 2024.

[15] Jieyi Li, Shuai Wang, Stevan Rudinac, and Anwar Osseyran. High-performance computing in healthcare:an automatic literature analysis perspective. *J. Big Data*, 11(1):61, 2024.

[16] Jieyi Li, Anwar Osseyran, Ruben Hekster, Stevan Rudinac, Valeriu Codreanu, and Damian Podareanu. Improving the speed and quality of cancer segmentation using lower resolution pathology images. *Multim. Tools Appl.*, 83(4):11999–12015, 2024.

[17] Inske Groenen, Stevan Rudinac, and Marcel Worring. Panorams: Automatic annotation for detecting objects in urban context. *IEEE Transactions on Multimedia*, 26:1281–1294, 2024.

[18] Jia-Hong Huang, Hongyi Zhu, Yixian Shen, Stevan Rudinac, Alessio M. Pacces, and Evangelos Kanoulas. A novel evaluation framework for image2text generation. In Clemencia Siro, Mohammad Aliannejadi, Hossein A. Rahmani, Nick Craswell, Charles L. A. Clarke, Guglielmo Faggioli, Bhaskar Mitra, Paul Thomas, and Emine Yilmaz, editors, *Proceedings of The First Workshop on Large Language Models for Evaluation in Information Retrieval (LLM4Eval 2024) co-located with 10th International Conference on Online Publishing (SIGIR 2024), Washington D.C., USA, July 18, 2024*, volume 3752 of *CEUR Workshop Proceedings*, pages 51–65. CEUR-WS.org, 2024.

[19] Omar Shahbaz Khan, Ujjwal Sharma, Hongyi Zhu, Stevan Rudinac, and Björn Pór Jónsson. Exquisitor at the lifelog search challenge 2024: Blending conversational search with user relevance feedback. In *Proceedings of the 7th Annual ACM Workshop on the Lifelog Search Challenge*, LSC '24, page 117–121, New York, NY, USA, 2024. Association for Computing Machinery.

[20] Hongyi Zhu, Jia-Hong Huang, Stevan Rudinac, and Evangelos Kanoulas. Enhancing interactive image retrieval with query rewriting using large language models and vision language models. In *Proceedings of the 2024 International Conference on Multimedia Retrieval*, ICMR '24, page 978–987, New York, NY, USA, 2024. Association for Computing Machinery.

[21] Stevan Rudinac, Alan Hanjalic, Cynthia Liem, Marcel Worring, Björn Pór Jónsson, Bei Liu, and Yoko Yamakata. *MultiMedia Modeling: 30th International Conference, MMM 2024, Amsterdam, The Netherlands, January 29 – February 2, 2024, Proceedings, Part I-V*. Springer Nature, 2024.

[22] Shuai Wang, Jiayi Shen, Athanasios Efthymiou, Stevan Rudinac, Monika Kackovic, Nachoem Wijnberg, and Marcel Worring. Prototype-enhanced hypergraph learning for heterogeneous information networks. In *MultiMedia Modeling*, pages 462–476, Cham, 2024. Springer Nature Switzerland.

[23] Omar Shahbaz Khan, Hongyi Zhu, Ujjwal Sharma, Evangelos Kanoulas, Stevan Rudinac, and Björn Pór Jónsson. Exquisitor at the video browser showdown 2024: Relevance feedback meets conversational search. In *MultiMedia Modeling*, pages 347–355, Cham, 2024. Springer Nature Switzerland.

[24] Ujjwal Sharma, Stevan Rudinac, Joris Demmers, Willemijn van Dolen, and Marcel Worring. Greenscreen: A multimodal dataset for detecting corporate greenwashing in the wild. In *MultiMedia Modeling*, Cham, 2024. Springer Nature Switzerland.

[25] Maarten Sukel, Stevan Rudinac, and Marcel Worring. Gigo, garbage in, garbage out: An urban garbage classification dataset. In *MultiMedia Modeling: 29th International Conference, MMM 2023, Bergen, Norway, January 9–12, 2023, Proceedings, Part I*, page 527–538, Berlin, Heidelberg, 2023. Springer-Verlag.

[26] Duc-Tien Dang-Nguyen, Cathal Gurrin, Martha Larson, Alan F Smeaton, Stevan Rudinac, Minh-Son Dao, Christoph Trattner, and Phoebe Chen. *MultiMedia Modeling: 29th International Conference, MMM 2023, Bergen, Norway, January 9–12, 2023, Proceedings, Part II*, volume 13834. Springer Nature, 2023.

[27] Mathias Glistrup, Stevan Rudinac, and Björn Pór Jónsson. Urban image geo-localization using open data on public spaces. In *Proceedings of the 19th International Conference on Content-Based Multimedia Indexing*, CBMI '22, page 50–56, New York, NY, USA, 2022. Association for Computing Machinery.

[28] Jonathan Bright, Nahema Marchal, Bharath Ganesh, and Stevan Rudinac. How Do Individuals in a Radical Echo Chamber React to Opposing Views? Evidence from a Content Analysis of Stormfront. *Human Communication Research*, 48(1):116–145, 11 2021.

[29] Athanasios Efthymiou, Stevan Rudinac, Monika Kackovic, Marcel Worring, and Nachoem Wijnberg. Graph neural networks for knowledge enhanced visual representation of paintings. In *Proceedings of the 29th ACM International Conference on Multimedia*, MM '21, page 3710–3719, New York, NY, USA, 2021. Association for Computing Machinery.

[30] Omar Shahbaz Khan, Björn Pór Jónsson, Jan Zahálka, Stevan Rudinac, and Marcel Worring. Impact of interaction strategies on user relevance feedback. In *Proceedings of the 2021 International Conference on Multimedia Retrieval*, ICMR '21, page 590–598, New York, NY, USA, 2021. Association for Computing Machinery.

[31] Omar Shahbaz Khan, Aaron Duane, Björn Pór Jónsson, Jan Zahálka, Stevan Rudinac, and Marcel Worring. Exquisitor at the lifelog search challenge 2021: Relationships between semantic classifiers. In *Proceedings of the 4th Annual on Lifelog Search Challenge*, LSC '21, page 3–6, New York, NY, USA, 2021. Association for Computing Machinery.

[32] Stevan Rudinac, Jenny Benois-Pineau, and Stephane Marchand-Maillet. Special issue on content-based multimedia indexing in the era of artificial intelligence. *Multimedia Tools and Applications*, pages 1–2, 2021.

[33] Stevan Rudinac, Alessandro Bozzon, Tat-Seng Chua, Suzanne Little, Daniel Gatica-Perez, and Kiyoharu Aizawa. Urbanmm'21: 1st international workshop on multimedia computing for urban data. In *Proceedings of the 29th ACM International Conference on Multimedia*, MM '21, page 5696–5697, New York, NY, USA, 2021. Association for Computing Machinery.

[34] Omar Shahbaz Khan, Björn Pór Jónsson, Mathias Larsen, Liam Poulsen, Dennis C. Koelma, Stevan Rudinac, Marcel Worring, and Jan Zahálka. Exquisitor at the video browser showdown 2021: Relationships between semantic classifiers. In Jakub Lokoč, Tomáš Skopal, Klaus Schoeffmann, Vasileios Mezaris, Xirong Li, Stefanos Vrochidis, and Ioannis Patras, editors, *MultiMedia Modeling*, pages 410–416, Cham, 2021. Springer International Publishing.

[35] Maarten Sukel, Stevan Rudinac, and Marcel Worring. Detecting urban issues with the object detection kit. In *Proceedings of the 28th ACM International Conference on Multimedia*, MM '20, page 4518–4520, New York, NY, USA, 2020. Association for Computing Machinery.

[36] Maarten Sukel, Stevan Rudinac, and Marcel Worring. Urban object detection kit: A system for collection and analysis of street-level imagery. In *Proceedings of the 2020 International Conference on Multimedia Retrieval*, ICMR '20, page 509–516, New York, NY, USA, 2020. Association for Computing Machinery.

[37] Omar Shahbaz Khan, Mathias Dybkjær Larsen, Liam Alex Sonto Poulsen, Björn Pór Jónsson, Jan Zahálka, Stevan Rudinac, Dennis Koelma, and Marcel Worring. Exquisitor at the lifelog search challenge 2020. In *Proceedings of the Third Annual Workshop on Lifelog Search Challenge*, LSC '20, page 19–22, New York, NY, USA, 2020. Association for Computing Machinery.

[38] Ujjwal Sharma, Stevan Rudinac, Marcel Worring, Joris Demmers, and Willemijn van Dolen. Semantic path-based learning for review volume prediction. *Advances in Information Retrieval: 42nd European Conference on IR Research*, ECIR 2020, Lisbon, Portugal, April 14–17, 2020.

[39] Omar Shahbaz Khan, Björn Pór Jónsson, Stevan Rudinac, Jan Zahálka, Hanna Ragnarsdóttir, Pórhildur Þorleiksdóttir, Gylfi Pór Guðmundsson, Laurent Amsaleg, and Marcel Worring. Interactive learning for multimedia at large. *Advances in Information Retrieval: 42nd European Conference on IR Research*, ECIR 2020, Lisbon, Portugal, April 14–17, 2020.

[40] Iva Gornishka, Stevan Rudinac, and Marcel Worring. Interactive search and exploration in discussion forums using multimodal embeddings. In Wen-Huang Cheng, Junmo Kim, Wei-Ta Chu, Peng Cui, Jung-Woo Choi, Min-Chun Hu, and Wesley De Neve, editors, *MultiMedia Modeling*, pages 388–399, Cham, 2020. Springer International Publishing.

[41] Björn Pór Jónsson, Omar Shahbaz Khan, Dennis C. Koelma, Stevan Rudinac, Marcel Worring, and Jan Zahálka. Exquisitor at the video browser showdown 2020. In Wen-Huang Cheng, Junmo Kim, Wei-Ta Chu, Peng Cui, Jung-Woo Choi, Min-Chun Hu, and Wesley De Neve, editors, *MultiMedia Modeling*, pages 796–802, Cham, 2020. Springer International Publishing.

[42] Devanshu Arya, Stevan Rudinac, and Marcel Worring. Hyperlearn: A distributed approach for representation learning in datasets with many modalities. In *Proceedings of the 27th ACM International Conference on Multimedia*, MM '19, pages 2245–2253, New York, NY, USA, 2019. ACM.

[43] Maarten Sukel, Stevan Rudinac, and Marcel Worring. Multimodal classification of urban micro-events. In *Proceedings of the 27th ACM International Conference on Multimedia*, MM '19, pages 1455–1463, New York, NY, USA, 2019. ACM.

[44] Hanna Ragnarsdóttir, Þórhildur Þorleiksdóttir, Omar Shahbaz Khan, Björn Pór Jónsson, Gylfi Pór Guðmundsson, Jan Zahálka, Stevan Rudinac, Laurent Amsaleg, and Marcel Worring. Exquisitor: Breaking the interaction barrier for exploration of 100 million images. In *Proceedings of the 27th ACM International Conference on Multimedia*, MM '19, pages 1029–1031, New York, NY, USA, 2019. ACM.

[45] Omar Shahbaz Khan, Björn Pór Jónsson, Jan Zahálka, Stevan Rudinac, and Marcel Worring. Exquisitor at the lifelog search challenge 2019. In *Proceedings of the ACM Workshop on Lifelog Search Challenge*, LSC '19, pages 7–11, New York, NY, USA, 2019. ACM.

[46] Devanshu Arya, Stevan Rudinac, and Marcel Worring. Predicting behavioural patterns in discussion forums using deep learning on hypergraphs. In *2019 International Conference on Content-Based Multimedia Indexing (CBMI)*, pages 1–6, Sep. 2019.

[47] Cathal Gurrin, Björn Pór Jónsson, Renaud Péteri, Stevan Rudinac, Stéphane Marchand-Maillet, Georges Quénot, Kevin McGuinness, Gylfi Pór Guðmundsson, Suzanne Little, Marie Katsurai, and Graham Healy, editors. *2019 International Conference on Content-Based Multimedia Indexing, CBMI 2019, Dublin, Ireland, September 4-6, 2019*. IEEE, 2019.

[48] Jan Zahálka, Stevan Rudinac, Björn Thór Jónsson, Dennis C. Koelma, and Marcel Worring. Blackthorn: Large-scale interactive multimodal learning. *IEEE Transactions on Multimedia*, 20(3):687–698, March 2018.

[49] Stevan Rudinac, Tat-Seng Chua, Nicolas Diaz-Ferreyra, Gerald Friedland, Tatjana Gornostaja, Benoit Huet, Rianne Kaptein, Krister Lindén, Marie-Francine Moens, Jaakko Peltonen, Miriam Redi, Markus Schedl, David A. Shamma, Alan Smeaton, and Lexing Xie. Rethinking summarization and storytelling for modern social multimedia. In *MultiMedia Modeling*, pages 632–644, Cham, 2018. Springer International Publishing.

[50] Stevan Rudinac, Iva Gornishka, and Marcel Worring. Multimodal classification of violent online political extremism content with graph convolutional networks. In *Proceedings of Thematic Workshops at ACM Multimedia Conference*, Mountain View, CA USA, 2017.

[51] Stevan Rudinac, Jan Zahálka, and Marcel Worring. Discovering geographic regions in the city using social multimedia and open data. In *MultiMedia Modeling: 23rd International Conference, MMM 2017, Reykjavík, Iceland, January 4-6, 2017.*, pages 148–159, Cham, 2017. Springer International Publishing.

[52] Joost Boonzajer Flaes, Stevan Rudinac, and Marcel Worring. What multimedia sentiment analysis says about city liveability. In *Advances in Information Retrieval: 38th European Conference on IR Research, ECIR 2016, Padua, Italy, March 20–23, 2016.*, pages 824–829, Cham, 2016. Springer International Publishing.

[53] Jorrit van den Berg, Stevan Rudinac, and Marcel Worring. Scenemash: Multimodal route summarization for city exploration. In *Advances in Information Retrieval: 38th European Conference on IR Research, ECIR 2016, Padua, Italy, March 20–23, 2016.*, pages 833–836, Cham, 2016. Springer International Publishing.

[54] Jan Zahálka, Stevan Rudinac, Björn Thór Jónsson, Dennis C. Koelma, and Marcel Worring. Interactive multimodal learning on 100 million images. In *Proceedings of the 2016 ACM International Conference on Multimedia Retrieval*, ICMR '16, pages 333–337, New York, NY, USA, 2016. ACM.

[55] Masoud Mazloom, Robert Rietveld, Stevan Rudinac, Marcel Worring, and Willemijn van Dolen. Multimodal popularity prediction of brand-related social media posts. In *Proceedings of the 2016 ACM Conference on Multimedia*, MM '16, pages 197–201, New York, NY, USA, 2016. ACM.

[56] Björn Thór Jónsson, Marcel Worring, Jan Zahálka, Stevan Rudinac, and Laurent Amsaleg. Ten research questions for scalable multimedia analytics. In *MultiMedia Modeling: 22nd International Conference, MMM 2016, Miami, FL, USA, January 4-6, 2016.*, pages 290–302, Cham, 2016. Springer International Publishing.

[57] Jan Zahálka, Stevan Rudinac, and Marcel Worring. Interactive multimodal learning for venue recommendation. *IEEE Transactions on Multimedia*, 17(12):2235–2244, Dec 2015.

[58] Jan Zahálka, Stevan Rudinac, and Marcel Worring. Analytic quality: Evaluation of performance and insight in multimedia collection analysis. In *Proceedings of the 23rd ACM International Conference on Multimedia*, MM '15, pages 231–240, New York, NY, USA, 2015. ACM.

[59] Judith Redi and Stevan Rudinac. Crowdmm 2015: Fourth international acm workshop on crowdsourcing for multimedia. In *Proceedings of the 23rd ACM International Conference on Multimedia*, MM '15, pages 1341–1342, New York, NY, USA, 2015. ACM.

[60] Stevan Rudinac and Marcel Worring. Making use of semantic concept detection for modelling human preferences in visual summarization. In *Proceedings of the 2014 International ACM Workshop on Crowdsourcing for Multimedia*, CrowdMM '14, pages 41–44, New York, NY, USA, 2014. ACM.

[61] Jan Zahálka, Stevan Rudinac, and Marcel Worring. New yorker melange: Interactive brew of personalized venue recommendations. In *Proceedings of the 22nd ACM International Conference on Multimedia*, MM '14, pages 205–208, New York, NY, USA, 2014. ACM.

[62] Stevan Rudinac. *Advancing the Relevance Criteria for Video Search and Visual Summarization*. Phd, Delft University of Technology, Delft, 05/2013 2013.

- [63] Stevan Rudinac, Alan Hanjalic, and Martha Larson. Generating visual summaries of geographic areas using community-contributed images. *IEEE Transactions on Multimedia*, 15(4):921–932, June 2013.
- [64] Stevan Rudinac, Martha Larson, and Alan Hanjalic. Learning crowdsourced user preferences for visual summarization of image collections. *IEEE Transactions on Multimedia*, 15(6):1231–1243, Oct 2013.
- [65] Stevan Rudinac, Martha Larson, and Alan Hanjalic. Leveraging visual concepts and query performance prediction for semantic-theme-based video retrieval. *International Journal of Multimedia Information Retrieval*, 1(4):263–280, 2012.
- [66] Martha Larson, Mohammad Soleymani, Pavel Serdyukov, Stevan Rudinac, Christian Wartena, Vanessa Murdock, Gerald Friedland, Roeland Ordelman, and Gareth J. F. Jones. Automatic tagging and geotagging in video collections and communities. In *Proceedings of the 1st ACM International Conference on Multimedia Retrieval*, ICMR ’11, pages 51:1–51:8, New York, NY, USA, 2011. ACM.
- [67] Stevan Rudinac, Alan Hanjalic, and Martha Larson. Finding representative and diverse community contributed images to create visual summaries of geographic areas. In *Proceedings of the 19th ACM International Conference on Multimedia*, MM ’11, pages 1109–1112, New York, NY, USA, 2011. ACM.
- [68] Stevan Rudinac, Martha Larson, and Alan Hanjalic. Tud-mir at mediaeval 2011 genre tagging task: Query expansion from a limited number of labeled videos. In *MediaEval Workshop*, 2011.
- [69] Stevan Rudinac, Martha Larson, and Alan Hanjalic. Exploiting result consistency to select query expansions for spoken content retrieval. In *Proceedings of the 32nd European Conference on Advances in Information Retrieval*, ECIR’2010, page 645–648, Berlin, Heidelberg, 2010. Springer-Verlag.
- [70] Stevan Rudinac, Martha Larson, and Alan Hanjalic. Exploiting noisy visual concept detection to improve spoken content based video retrieval. In *Proceedings of the International Conference on Multimedia*, MM ’10, pages 727–730, New York, NY, USA, 2010. ACM.
- [71] Stevan Rudinac, Martha Larson, and Alan Hanjalic. Visual concept-based selection of query expansions for spoken content retrieval. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR ’10, pages 891–892, New York, NY, USA, 2010. ACM.