

LORIS BAZZANI

PRINCIPAL AI SCIENTIST AT AMAZON AND MUSICIAN

loris.bazzani@gmail.com <http://lorisbaz.github.io>

<http://scholar.google.com/citations?user=1cdNGL4AAAAJ&hl=en>

CURRENT POSITION

Principal AI Scientist at Amazon

2016-now

- Passionate about multimodal (generative) AI, that connect images/videos, language and different types of user interactions (e.g., LMMs and stable diffusion)
- Capacity to reason about a wide breath of science and products, and ability to switch to science depth mode at the right time
- Ability to design and prototype complex AI models and demonstrate their practical value to leadership
- Experience in transferring technology into a shipping product at Amazon scale. Examples of successful projects in production I contributed to:
 - Product recommendation widgets that leverage multimodal similarity with diversification, trained at scale for fashion and home decor. Accessible by millions of customer in the detail page of products and to Alexa-enabled devices
 - Virtual try-on experiences for make-up usable via the Amazon mobile App
 - Rapid Recap an action recognition-based tool to catch up on plays during live football games in Prime Video
 - Detection of nearly-similar videos and maturity rating to reduce manual tedious audit tasks in Prime Video
- Responsible of defining the strategy and science direction of a team of scientists. Experience in establishing research and productization roadmaps
- Collaborate and partner with business and engineering teams across the company
- Contribute in creating a centralized science team from 4 to 15 scientists
- Diversity, inclusion and accessibility ally
- Ability to present my work to non-technical audience and discuss with senior leadership, including VPs and SVPs
- Mentorship experience of interns (16), scientists, engineers and product managers

ACADEMIC EXPERIENCE & EDUCATION

Dartmouth College

2014–2015

Postdoc at the Visual Learning group led by L. Torresani

- Research on object localization and detection in images, deep learning and salience prediction for videos

Italian Institute of Technology

2012–2013

Postdoc at the Pattern Analysis and Computer Vision group led by V. Murino

- Research on social interaction analysis in videos using dynamic Bayesian nets and Bayesian non-parametric models, and on kernel-based object recognition and person re-identification

University of Verona

2009–2012

Ph.D. in Computer Science with focus on Computer Vision and Machine Learning

- PhD Thesis: Beyond multi-target tracking: statistical pattern analysis of people and groups
- Supervisors: V. Murino and M. Cristani. Reviewers: A. Del Bimbo and R. T. Collins

University of British Columbia

2010

Research Intern at the Laboratory for Computational Intelligence led by N. de Freitas

- Research on restricted Boltzmann machines, deep learning and attentional models

University of Verona

2006–2008

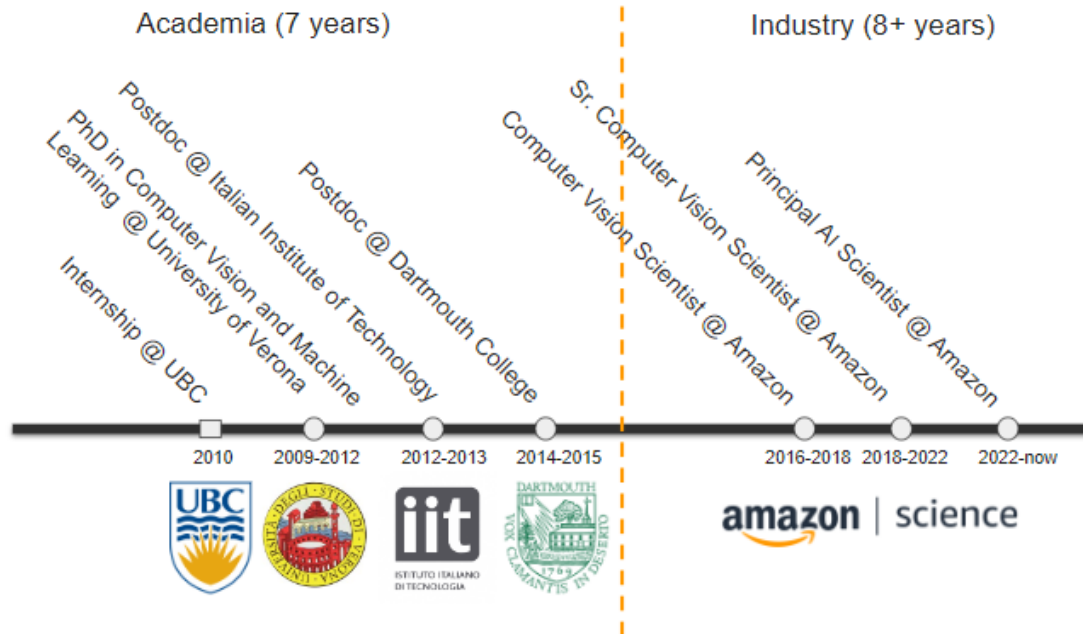
M.S. in Intelligent and Multimedia Systems. Thesis topics: particle filtering and multi-person tracking. Summa cum laude and top of the class

University of Verona

2003–2006

B.S. in Information Technology. Thesis topics: machine learning for medical imaging

EXPERIENCE



RESEARCH, PUBLICATIONS & PATENTS

- A. Cao, M. Jaritz, M. Guillaumin, R. de Charette, L. Bazzani. LatteCLIP: Unsupervised CLIP Fine-Tuning via LMM-Synthetic Texts, WACV 2025.
- X. Yang, Y. Zuo, S. Ramasinghe, L. Bazzani, G. Avraham, A. van den Hengel. ViewFusion: Towards Multi-View Consistency via Interpolated Denoising, CVPR 2024.
- R. Bodur, E. Gundogdu, B. Bhattarai, T-K Kim, M. Donoser, L. Bazzani. iEdit: Localised Text-guided Image Editing with Weak Supervision, In CVPR Workshops, 2024.
- US Patent 11,720,942: Interactive retrieval using visual semantic matching, 2023.
- US Patent 11,809,520: Localized visual similarity, 2023.
- US Patent 11,829,445: Attribute-based interactive product recommendations, 2023.
- M. Xu, E. Gundogdu, M. Lapin, B. Ghanem, M. Donoser, L. Bazzani. Contrastive Language-Action Pre-training for Temporal Localization, Arxiv, 2022.
- US Patent 11,416,910: Visual blending of content, 2022.
- US Patent 11,361,212: Machine learning system to score alt-text in image data, 2022.
- Y. Hou, E. Vig, M. Donoser, L. Bazzani. Learning Attribute-driven Disentangled Representations for Interactive Fashion Retrieval, ICCV, 2021.
- A. Salvador, E. Gundogdu, L. Bazzani, M. Donoser. Revamping Cross-Modal Recipe Retrieval with Hierarchical Transformers and Self-supervised Learning, CVPR, 2021.
- A. D’Innocente, N. Garg, Y. Zhang, L. Bazzani, M. Donoser. Localized Triplet Loss for Fine-Grained Fashion Image Retrieval. CVPR Workshops, 2021.
- Y. Chen, and L. Bazzani. Learning Joint Visual Semantic Matching Embeddings for Language-guided Retrieval, ECCV, 2020.
- Y. Chen, S. Gong, and L. Bazzani. Image Search with Text Feedback by Visiolinguistic Attention Learning, CVPR, 2020.
- US Patent 10,643,074: Automated video ratings, 2020.
- L. Bazzani, T. Domhan, and F. Hieber. Image Captioning as Neural Machine Translation Task in SOCKEYE, Arxiv, 2018.
- L. Bazzani, H. Larochelle, and L. Torresani, Recurrent Mixture Density Network for Spatiotemporal Visual Attention, ICLR, 2017.

- S. Vascon, and L. Bazzani, Group Detection and Tracking using Sociological Features, In Group and Crowd Behavior for Computer Vision, 2017.
- M. San Biagio, H. Q. Minh, L. Bazzani, V. Murino. Approximate Log-Hilbert-Schmidt distances between covariance operators for image classification, [CVPR](#), 2016.
- L. Bazzani, A. Bergamo, D. Anguelov, L. Torresani. Self-taught object localization with deep networks, WACV, 2016.
- H. Q. Minh, L. Bazzani, V. Murino, A unifying framework in vector-valued reproducing kernel Hilbert spaces for manifold regularization and co-regularized multi-view learning, [JMLR](#), 2016.
- H. Q. Minh, M. San Biagio, L. Bazzani, and V. Murino, Kernel Methods on Approximate Infinite-Dimensional Covariance Operators for Image Classification, Arxiv, 2016.
- L. Bazzani, M. Zanotto, M. Cristani, V. Murino, Joint individual-group modeling for tracking, [PAMI](#), 2015.
- L. Bazzani, M. Cristani, and V. Murino. SDALF: Modeling human appearance with symmetry-driven accumulation of local features, Person Re-identification, 2014.
- M. San Biagio, L. Bazzani, M. Cristani, V. Murino, Weighted bag of visual words for object recognition, ICIP, 2014.
- H. Q. Minh, L. Bazzani, V. Murino, A unifying framework for vector-valued manifold regularization and multi-view learning, [ICML](#), 2013.
- D. Figueira, L. Bazzani, H.Q. Minh, M. Cristani, A. Bernardino, and V. Murino, Semi-supervised multi-feature learning for person re-identification, AVSS, 2013.
- P. Salvagnini, L. Bazzani, M. Cristani, and V. Murino, Person re-identification with a PTZ camera: an introductory study, ICIP, 2013.
- L. Bazzani, M. Cristani, and V. Murino. Symmetry-driven accumulation of local features for human characterization and re-identification, [CVIU](#), 2013.
- L. Bazzani, D. Tosato, M. Cristani, M. Farenzena, G. Paggetti, G. Menegaz, V. Murino, Social interactions by visual focus of attention in a three-dimensional environment, Expert Systems, 2013.
- L. Bazzani, V. Murino, and M. Cristani, Decentralized particle filter for joint individual-group tracking, [CVPR](#), 2012.
- M. Zanotto, L. Bazzani, M. Cristani, and V. Murino, Online bayesian non-parametrics for social group detection, [BMVC](#), 2012.
- L. Bazzani, M. Cristani, G. Paggetti, D. Tosato, G. Menegaz, V. Murino, Analyzing groups: a social signaling perspective, Video Analytics for Business Intelligence, 2012.
- B. I. Barbosa, M. Cristani, A. Del Bue, L. Bazzani, and V. Murino, Re-identification with RGB-D sensors, In 1st International Workshop on Re-Identification, 2012.
- L. Bazzani, M. Cristani, A. Perina, V. Murino, Multiple-shot person re-identification by chromatic and epitomic analyses, PRL, 2012.
- M. Denil, L. Bazzani, H. Larochelle, and N. de Freitas, Learning where to attend with deep architectures for image tracking, [Neural Computation](#), 2012.
- L. Bazzani, N. de Freitas, H. Larochelle, V. Murino, J. Ting, Learning attentional policies for object tracking and recognition in video with deep networks, [ICML](#), 2011.
- D. S. Cheng, M. Cristani, M. Stoppa, L. Bazzani, V. Murino, Custom pictorial structures for re-identification, [BMVC](#), 2011.
- M. Cristani, L. Bazzani, G. Paggetti, A. Fossati, A. Del Bue, D. Tosato, G. Menegaz, V. Murino, Social interaction discovery by statistical analysis of F-formations, [BMVC](#), 2011.
- M. Cristani, G. Paggetti, A. Vinciarelli, L. Bazzani, G. Menegaz, V. Murino, Towards computational proxemics: Inferring social relations from interpersonal distances, SocialCom, 2011.
- L. Bazzani, M. Cristani, A. Perina, M. Farenzena, and V. Murino, Multiple-shot person re-identification by HPE signature, ICPR, 2010.
- M. Farenzena, L. Bazzani, A. Perina, V. Murino, and M. Cristani Person re-identification by symmetry-driven accumulation of local features, [CVPR](#), 2010.

- L. Bazzani, N. de Freitas, J. Ting, Learning attentional mechanisms for simultaneous object tracking and recognition with deep networks, Workshop on Deep Learning and Unsupervised Feature Learning, 2010.
- L. Bazzani, M. Cristani, and V. Murino, Collaborative particles filters for group tracking, ICIP, 2010.

Blue color = top-tier computer vision/machine learning conferences or journals.

CODING & TOOLS

- Python, PyTorch, HuggingFace, git, LATEX, bash scripting, LLM prompting, various AWS services
- Passionate about and participated to several internal Hackathons (at least 1 per year)
- In past, I worked with MATLAB, Lua/Torch7, JavaScript, Qualtrics, Alexa skill devel tools

AWARDS & GRANTS

- **Amazon Accessibility Innovation Award**, internal award given to individuals who created experiences to improve accessibility, 2021.
- **Amazon Robotics Challenge Award**, internal award given to the team who won the robotics challenge, 2018.
- **IBM Best Student Paper Award**, track: Computer Vision at International Conference on Patter Recognition ICPR, 2010
- **Scholarship** from University of Verona that supported my Ph.D. from Jan. 2009 to Dec. 2011
- **Scholarship** from EU-Project FP7 SAMURAI, grant FP7-SEC-2007-01 No. 217899, that contributed to support my Ph.D. from Jan. 2009 to Dec. 2011
- **Travel grant** from University of British Columbia to attend Neural Information Processing Systems, 2010
- **Travel grant** from International Machine Learning Society to attend the International Conference on Machine Learning, 2011

AI COMMUNITY EFFORTS

- Co-organizer of the workshop on Computer Vision for Fashion, Art, and Design @ CVPR
- Co-organizer of the workshop on 3D Vision and Modeling Challenges in eCommerce @ ICCV
- Co-organizer and program chair of several AI conferences and shopping science summits internal at Amazon
- Active reviewer for CVPR, ICCV, ECCV, NeurIPS, ICML, ICLR, IEEE TPAMI, IJCV, IEEE Transactions on Multimedia, IEEE Transactions on Image Processing

TEACHING & PRESENTATIONS

Presentations

Delivered 50+ technical presentations at AI conferences and internal Amazon events to audiences of up to 1000 people

Dartmouth College

Machine Learning. Assistant of undergrad course by L. Torresani

Winter 2015

Italian Institute of Technology

An Introduction to Machine Learning and Computer Vision for PhD students

Winter 2013

University of Modena and Reggio Emilia

Person Re-identification at the Short Spring School in Surveillance for PhD students

Spring 2011

MY MENTORS

I was fortunate enough to work with and learn from my awesome mentors during my career so far. Their advice and feedback helped me growing in a way I could never imagine. Thank you all!

- Anton van den Hengel, Larry Davis, Raffay Hamid, Josh Miele, Werner Trobin: science direction and guidance at Amazon
- Michael Donoser and Rajiv Chopra: leadership and management guidance at Amazon

- Lorenzo Torresani and Hugo Larochelle at Dartmouth College
- Vittorio Murino and Marco Cristani at University of Verona and at the Italian Institute of Technology
- Nando de Freitas at the University of British Columbia

WHAT ELSE ABOUT ME

You might want to know other traits of me, not just my technical skills and achievements, right? Here are some:

- Growth mindset and value learning new skills, expanding my knowledge and improving different abilities in all aspects of my life
- Cultivate cross-cultural communication skills through living and working in diverse environments
- Nurture my creative side as musician, playing keyboards and piano since a young age. I played in several local bands and have one now. I'm currently learning guitar and singing
- Invest on taking care of the mind, emotions and my mental health via meditation, introspection and other tools
- High level of self-awareness. I'm keen to admit my mistakes when others or I recognize them, and learn from them
- My favourite books I've recently read: the power of vulnerability by B. Brown; the myth of normal by G. Matè and D. Matè; the five invitations by F. Ostaseski
- My big 5: openness 90, conscientiousness 88, extraversion 40, agreeableness 46 and neuroticism 29
- Fluently speak Italian, English and Venetian (yes, it's a language and quite different from Italian). I'm slowly learning German too

What colleagues say about me:

- "You cultivate an exceptional research culture where creativity and science can thrive. Through your guidance, teams uncover innovative paths forward by exploring new directions and questioning norms."
- "He is a top tier scientist that can handle large technical complexity. His scientific depth, breadth and technical skills are top notch."
- "One of his super powers is to give honest and constructive feedback, without compromising for social cohesion. He voices concerns when a minority of the group might also be impacted."
- "He is that rare combination between a big thinker and someone who's in touch with the technical details."

Some of the growth areas I'm focusing on: foster more autonomy in others via delegation and improve my strategical thinking by leveraging more on intuition and influence of others.