

BigScience project

Working Groups

Collaborative tasks

The workshop involves a set of **collaborative tasks** conducted along the year and aimed at creating, sharing and evaluating a large multilingual dataset and a large language model as tools for research. The collaborative tasks are quite large since they involve several millions GPU hours on a supercomputer.

The goal of the collaborative tasks is in particular to collaboratively create and share several **artifacts** as tools for the research community:

- A **Multilingual Dataset** for Research (possibly full-open-access, possibly behind an authentication access portal to follow GdPR and data privacy legislations => to be defined by the organization committee in charge of this artifact)
- A **large Language Model** for Research (possibly full-open-access, possibly behind a researcher's authentication access portal like the dataset => to be defined by the organization committee in charge of this artifact)
- **Software tools** created to complete the shared tasks, e.g. dataset filtering, model off-loading, etc (under very permissive license, e.g. Apache2 => to be by the organization committee in charge of this artifact)
- Extensive **documentation** of the evaluation results, protocols, approach, datasets, and tools developed in the course of the project for future reconduction

Table of Content

1. [Overview: how to participate - what roles](#)
2. [How to join, co-chair or propose a Working Group](#)
3. [More details on the role of the co-chairs](#)
4. [Working Groups already created](#) (Note that Chairs are evolving)
 - 4.1. [Workshop organization](#) - Chairs: Matthias Gallé, Angela Fan, Suzana Ilic
 - 4.2. [Dataset creation](#) - Area-Chair: Yacine Jernite
 - 4.2.1. [Data Governance and Archival Strategies](#) - Chairs: Yacine Jernite, Huu Nguyen, Meg Mitchell

- 4.2.2. [Legal and Ethical Scholarship](#) - Chairs: Carlos Muñoz Ferrandis, Giada Pistilli, Somaieh Nikpoor and David Lansky
- 4.2.3. [Privacy](#) - Chairs: Anima Shukla, Suhas Pai, Huu Nguyen, HESSIE Jones, Jenny Chim
- 4.2.4. [Data Sourcing and Representativeness](#) - Chairs: Angie McMillan-Major, Pedro Ortiz Suárez, Zeerak Waseem
- 4.2.5. [Data Tooling: from Data Sources to Training Dataset](#) - Chairs: Albert Villanova, Sampo Pyysalo, Filip Ginter
- 4.3. [Modeling approach](#) - Area-Chairs: Victor Sanh, Stephen Bach
 - 4.3.1. [Modeling – Architecture and Scaling](#) - Chairs: Julien Launay, Iz Beltagy
 - 4.3.2. [Modeling – Metadata](#) - Chairs: Timo Schick, Nora Kassner
 - 4.3.3. [Modeling – Multilinguality](#) - Chairs: Vassilina Nikoulina, Hady Elsahar
 - 4.3.4. [Modeling – Prompt engineering](#) - Chairs: Victor Sanh, Stephen Bach
 - 4.3.5. [Modeling – Retrieve](#) - Chair: Nils Reimers
- 4.4. [Tokenization](#) - Chairs: Benoît Sagot, Samson Tan
- 4.5. [Evaluation](#) - Chairs: Aurélie Névéol, Dan Garrette, Ellie Pavlick, Marine Carpuat, Thomas Scialom, Sebastian Gehrmann, Verena Rieser
 - 4.5.1. [Evaluation - Extrinsic Evaluation](#) - Chairs: Thomas Scialom, Verena Rieser
 - 4.5.2. [Evaluation - Intrinsic Evaluation](#) - Chairs: Sebastian Gehrmann
 - 4.5.3. [Evaluation - Few-Shot Generalization](#) - Chairs: Ellie Pavlick
 - 4.5.4. [Evaluation - Bias, Fairness, and Social Impact](#) - Chairs: Aurelie Neveol
 - 4.5.5. [Evaluation - Multilinguality](#) - Chairs: Dan Garrette & Marine Carpuat
- 4.6. [Interpretability and interaction/visualization](#) - Chair: Hendrik Strobelt, Ryan Teehan
- 4.7. [Engineering/scaling](#) - Chairs: Yozh (jheuristic), Max Ryabinin, Minjia Zhang
- 4.8. [Collaborations with education and research communities](#) - Chairs: Dieuwke Hupkes, François Yvon
- 4.9. [Carbon footprint](#) - Chairs: Anne-Laure Ligozat, Canwen Xu, Sasha Luccioni, Sylvain Viguier
- 4.10. [\[Meta-Group\] Social Impact Across Groups](#) - Chair: Karen Fort, Michael McKenna, Margot Mieskes
- 4.11. [Sharing and accessibility of model and dataset](#) - Chair: Thomas Wolf, Valentin Danchev, Danish Contractor and Aaron Gokaslan
- 4.12. [Challenges in the biomedical domain](#) - Chairs: Jason Fries, Debajyoti Datta, Natasha Seelam
- 4.13. [Language Modeling for Mathematics and Theorem Proving](#) - Chairs: Jelle Piepenbrock, Bartosz Piotrowski

5. [Proposal for New Working Groups](#)

5.1. [Template to add a new working group](#)

Overview: how to participate - what roles

BigScience is a research workshop and is open to researchers, i.e. people affiliated with a research organization (in academia or industry) and whose day work is for instance (at least in part) to publish papers in peer-review venues, as well as to people whose technical and professional expertise bears relevance to the social aspects of the project.

The organization of the research workshop is not a good place to learn about AI or find mentors as we do not have an internal mentoring program. It's also not per-se a community for general AI or AGI discussions.

If you are interested in learning about AI/ML or already working on Machine Learning, AI and NLP side projects, note that we plan to organize events for the AI/ML enthusiast community and the general public. You can fill out the following [Google Form](#) and join the mailing list to be informed about these upcoming public events

You can generally participate in the project as:

- **Advisor (*Everyone is advisor by default when joining*)**
 - *role*: give general scientific/organization advices
 - *time commitment*: light - reading a newsletter every 2 weeks - giving feedback/advices
- **Participant in a Working Group (*Add your name below*)**
 - *role*: joining one of the working groups of the OC (see list below): advising/designing/building the collaborative task (building the dataset/model/tools) or advising/designing/organizing the live events
 - *time commitment*: medium - depend on the chosen task (see details of the working groups below)
- **Chair/co-chair of a Working Group (*Add your name below*)**
 - *role*: the chair(s) is supposed to provide at least the minimal amount of work necessary for having a very bare-bone version of the task. If other members are active in the WG, the chair(s) can mostly coordinate the effort and organise the decision process.
 - *time commitment*: more significant - also depend on the chosen WG
- **Workshop attendant** joining the collaborative task or events (***To be defined later***)

- *role*: participating in the collaborative task in a guided way following guidelines setup by the OC (helping build the dataset, helping build the tools)
- *time commitment*: up to the attendant - following guidelines of the OC/WG

How to join, co-chair or propose a Working Group

If you want to propose a new group, ask to join a Working Group or ask to chair/co-chair a Working Group, well, first of all, this is great, *thanks a lot!*

You can follow the following steps:

- **propose** a new working group
Copy the template at the end of this document in the “Proposal for New Working Groups” section. Feel free to propose yourself as a Chair or just propose a Working Groups idea. You should fill a-minima the “Scope”. (add in “suggestion” mode so we receive a notification that you’ve added it)
- **participate** in a working group
Add your name in the “Current member” section of the working group (add in “suggestion” mode so we receive a notification that you’ve added it)
- **be chair or co-co-chair** an existing or a new working group
Add your name in the “Current chairs” section of the working group (add in “suggestion” mode so we receive a notification that you’ve added it)
Working groups can have 3-4 chaires without problems.

More details on the role of the co-chairs

Each working group has at least one chair (ideally more than 1 person). Here are some details on the chair role and selection.

The chairs role:

- chairs are supposed to coordinate the effort in each working group and organise the decision process while letting working group members feel welcome to raise ideas and propose things
- if no-one is active in the working group, chairs are supposed to do at least the minimal amount of work necessary for having a bare-bone version of the task (aka a minimal "fallback", e.g. gathering Common Crawl for the dataset working groups)

- chairs are in charge of making regular summary of the work to share with all and participate in all-chairs-sync meetings
- chairs are in charge of maintaining the on-boarding resource to have people join the working group
- chairs are responsible to ensure respectful and open discussions in their WG. As such, they should in particular ensure that the [code of conduct](#) is respected.

For the **selection process**, we are trying to have something quite bottom-up keeping in mind that chair's role are quite essential to the success of the project so there is a validation step and we (Thomas Wolf in particular) may ask people to not join as chairs or step down if there seems to be a discrepancy between the requirements of the role and the current investment of a candidate. Here is the general selection process

- generally asking for self-propositions for chairs or co-chairs in the current document,
- if no one show up the organization working group will reach out to specific people with related skill and interest in participating,
- when additional people want to be chairs, the current chairs should have a voice to agree or not to the additional chairs,
- when a decision needs to be made to select a chair, the Workshop Organization working group will vote a decision,
- in all this process goodwill is assumed and there should be an effort to listen to all arguments but the pragmatic success of the project is also considered more important than specific people interest (in other words "wanting to be a chair does not guarantee that you will be a chair and there will be a certain degree of arbitrariness", I'm sorry in advance).

Working Groups already created

1. Workshop organization

Scope

This working group is in charge with generally defining how the collaboration could be structured (as a workshop) and organizing the live events and activities not directly related to the collaborative tasks during the one year duration of the workshop as well as communication tools (newsletters, forum, etc).



Resources - onboarding - documentation

Folder on the google drive: [Organization](#)

This folder includes past newsletters summarizing the evolution of the project as well as general organizational documents and documents related to the live events and other activities not directly related to the collaborative tasks.

Current members:

Matthias Gallé, Angela Fan, Y-Lan Boureau, Hendrik Strobelt, Verena Rieser, Thomas Wolf, Victor Sanh, Yacine Jernite, Suzana Ilic, Sebastian Gehrmann, Ellie Pavlick, Shubham Agarwal, Meheresh Masanpally, Christopher Akiki, Alya Karim, Seong-Young Her

Current chair(s)

Angela Fan, Matthias Gallé and Suzana Ilic

Current status

Currently planning the tools (website + forum), organizing the first kick-off live event, discussing the general workshop setup, inviting members to the workshop, discussing with (public and private) compute providers

2. Dataset creation

Scope

The Dataset Creation effort is concerned with the creation of a large multilingual dataset. It is further split into 3 distinct working groups:

- [Data Governance and Archival Strategies](#)
- [Data Sourcing and Representativeness](#)
- [Data Tooling: from Data Sources to Training Dataset](#)

Resources - onboarding - documentation

- Folder on the google drive: [Dataset documents](#)
- Entry document: [BigScience Data Section](#)



Current Area Chair(s)

Yacine Jernite

2.1. Data Governance and Archival Strategies

Scope

This working group is tasked with exploring questions about how data is gathered and managed, choosing the right metadata and indexing and documentation structure, and developing protocols to ensure that data is used while respecting the rights and for the benefit of the data subjects.

- Defining a management and ownership structure for the dataset
- Scoping out legal concerns and societal impact of data choices
- Providing recommendations for how the above should affect the indexing, metadata, sourcing, and data maintenance infrastructure

Resources - onboarding - documentation

Entry document: [Dataset Org: Data Governance and Archival Strategies](#)

Current members

Karën Fort, Guillaume Klein, Samson Tan, Jordi Armengol-Estapé, Yacine Jernite, Félicien Vallet, Luca Soldaini, Alison Callahan, Madhulika Srikumar, Stella Biderman, Alicia Nobles, Anouk Ruhaak, Zeerak Waseem, Huu Nguyen, Jesse Dodge, Gérard Dupont, Shanya Sharma, Maraim Masoud, Benoît Sagot, Marta Villegas, Nicolas Hervé, Pedro Ortiz Suárez, Filip Ginter, Sebastian Nagel, Trieu Le, Michael McKenna, Jade Abbott, Kyle Lo, Valentin Danchev, Ebenge Usip, Manan Dey, Zaid Alyafeai, Debajyoti Datta, Isaac Johnson, Martin Gerlach, Rose E. Wang, Charles Welch, Suhas Pai, Daniel van Strien, Long Phan, Binh Nguyen, Duong A. Nguyen, Edward Tan, Hieu Tran, Thanh Le, Trieu Le, Habiba Dahmani, Chenghao Mou, Olivier Nguyen, Adela Ljajić, Amit Alfassy, Jonathan Chang, Dragomir Radev, Jenny Chim, Rishi Bommasani, Shayne Longpre, Ian yu, Ryan Hao, Mike Qiu, An Ran, Yoyo Yang, Zheng Xin Yong, Christopher Akiki, Keith Stevens, Simon Wallace, Alya Karim

Current chair(s)

Yacine Jernite, Meg Mitchell



Current status

Scoping out further collaborative tasks within the group.

2.1.2 Legal and Ethical Scholarship

Scope

This sub working group discusses legal topics related to data governance and feeds into the larger data governance group.

Resources - onboarding - documentation

Entry document: [Diagnostics of Legal Landscape for BigScience](#)

Current members

Anna Tordjmann, Carlos Muñoz Ferrandis, David Lansky, Huu Nguyen, Meg Mitchell, Michael McKenna, Sharanya, Yacine Jernite, Kevin Jordan, Somaieh Nikpoor, Sergi Gálvez Duran, Marina Valls, Michael Weinberg, Danda Zhao, Anima Shukla, Madhulika Srikumar, Aaron Gokaslan, Tara Davis, Suhas Pai, Giada Pistilli, Jesse Dodge, Kyle Lo, Peter Henderson, Ramya Chandrasekhar, Maria A. Castillo, Ishani Dash, Minna Liu, Helena U. Vrabec, Luisa Shinzato, Caio Brito, Olivier Nguyen, Tomoya Kainuma, Yingxin Xu, Kenneth Heafield, Danish Contractor, Long Phan, Binh Nguyen, Duong A. Nguyen, Edward Tan, Hieu Tran, Thanh Le, Trieu Le, Chenghao Mou, Olivier Nguyen, Bec Johnson, Jenny Chim, Maiko Takeuchi, Sinee Sang-aroonsiri, Yu Xu, Daniel Molano, Rodrigo Canalli, Rishi Bommasani, Eduardo González Ponferrada, Javier de la Rosa, Manuel Romero, Maria Grandury, and Paulo Villegas, Hessie Jones, Pawan Sasanka Ammanamnchi, Jörg Frohberg, Imane Bello, Pete Smith, Simon Wallace

Current chair(s)

Carlos Muñoz Ferrandis, Giada Pistilli, Somaieh Nikpoor and David Lansky

2.1.3 Privacy

Scope

This sub working group researches technology, policy, ethics and legal issues related to privacy protection in datasets, modeling and usage. This WG will interface with other WGs focusing on rights protections and social impacts, including Legal Scholarship, Data Tooling, Data Governance, Model Governance, Social Impacts, and potentially Evaluation. This WG will also run the PII Hackathon and perform PII remediation on the final data used to train the final BigScience model.

Resources - onboarding - documentation

Entry document: [BigScience Privacy Protection](#)

Current members

Anna Rogers, Colin Raffel, David Adelani, Elizabeth Keleshian, Hady Elsahar, Huu Nguyen, Jason Fries, Jesse Dodge, MaraimM, Matthias Galle, Meg Mitchell, Nora Kassner, Oliver Nguyen, Robert Monarch, Sasha Luccioni, Shanya Sharma, Somaieh Nikpoor, Stella Biderman, Tal Perry, Victor Sanh, Yacine Jernite, Zeerak Waseem, Anima Shukla, Amir Feizpour, Suhas Pai, Olivier Nguyen, Javier de la Rosa, Manuel Romero, Francesco De Toni, Long Phan, Binh Nguyen, Duong A. Nguyen, Edward Tan, Hieu Tran, Thanh Le, Trieu Le, Chenghao Mou, Olivier Nguyen, Ian Yu, Jenny Chim, Hessie Jones, Julio Bonis Sanz, Jason Fries, Rishi Bommasani, Christopher McMaster, Minna Liu, Helena U. Vrabec, Pawan Sasanka Ammanamanchi, Nishant Subramani (he/him), Omri Mendels, Joseph Allen, Shamik Bose, Jörg Frohberg, Imane Bello, Daniel Huynh

Current chair(s)

Anima Shukla, Suhas Pai, Huu Nguyen, Hessie Jones, Jenny Chim

2.2. Data Sourcing and Representativeness

Scope

This working group is responsible for helping define language choices and local and global representativeness criteria, analysing the diversity of existing text sources for each region in terms of social contexts represented, and finding diverse sources of text to meet these criteria, including both online and offline text in all available media.

- Defining frameworks for analysing representativeness and diversity across regions
- Exploring different modes of data collection from web crawling to participatory methods and collaboration with existing data organizations
- Choosing languages, defining language varieties, and identifying relevant regions for each of those
- Identifying diverse text sources for each region and language variety

Resources - onboarding - documentation

Entry document: [Dataset Org: Data Sourcing and Representativeness](#)

Current members

Karën Fort, Sam Bowman, Halil Akin, Caiming Xiong, Guillaume Klein, Samson Tan, Myle Ott, Philippe Muller, Ruiqi Zhong, Luke Zettlemoyer, Yacine Jernite, Wietse de Vries, Max Ryabinin, Antoine Neuraz, Tsvetomila Mihaylova, Hady Elsahar, Manan Dey, Shanya Sharma, Minh Quang Pham, Jin Koay, Ari Jankelowitz, Elizabeth Keleshian, Shiyue Zhang, Evan Dufraisse, Edoardo M. Ponti, Han Wang, Ona de Gibert Bonet, Zaid Alyafeai, Md Rabiul Awal, Kaustubh Dhole, Jonathan Chang, Maximin Coavoux, Adrian Popescu, Maraim Masoud, Ben Peters, Tasnim Mohiuddin, Rabin Banjade, Vinay Uday Prabhu, Aitor Soroa, Trishala Neeraj, Luca Soldaini, Rodrigo Wilkens, Canwen Xu, Sheng Shen, Michael McKenna, Rishi Bommasani, Patrick Drouin, Fredrik Olsson, Sadid A. Hasan, Francesco De Toni, Huu Nguyen, Laurent Besacier, Nicolas Hervé, Ludovic Tanguy, Benoît Sagot, Salomey Osei, Alham Fikri Aji, Filip Ginter, Sampo Pyysalo, Gérard Dupont, Aakanksha Naik, Olivier Nguyen, Trieu Le, Emily Reif, Tolga Bolukbasi, Fabrice Zepfack, Max Huang, Mustafa Ghaleb, Aaron Gokaslan, Leo Gao, Olanrewaju Samuel, Taewoon Kim, Chris Emezue, Suhas Pai, Daniel van Strien, Long Phan, Binh Nguyen, Duong A. Nguyen, Edward Tan, Hieu Tran, Thanh Le, Trieu Le, Max Bartolo, Chenghao Mou, Olivier Nguyen, Birger Moëll, Adela Ljajić, Levkovizh Efrat, Ariel Kreisberg Nitzav, Eyal Bar Natan, Stanislav Silberberg, Khalid Almubarak, Haoming Jiang, Rishi Bommasani, Jörg Froberg, Shayne Longpre, Aashiq Muhamed, Keith Stevens, Samuel Cahyawijaya, Yada Pruksachatkun

Current chair(s)

Angie McMillan-Major, Pedro Ortiz Suárez, Zeerak Waseem

Current status

Scoping out further collaborative tasks within the group.

2.3. Data Tooling: from Data Sources to Training Dataset

Scope

This working group will develop tools to gather text from the identified sources and process it to be both easy to use at training time and respectful of the data subjects' rights. This includes tools for crawling, automatic PI detection and de-identification, documentation of web content, efficient web text formats that retain the website and page structure, tools for extracting text from audio or pdf files, as well as infrastructure for securely maintaining and dispensing the data for training.

Resources - onboarding - documentation

Entry document: [BigScience Data Tooling Working Group](#) (see also [Dataset Org: from Data Sources to Training Dataset](#))

Current members

Benoit Dal Ferro, Xavier Tannier, Halil Akin, Pasquale Minervini, Antoine Simoulin, Shubham Agarwal, Yozh, Yacine Jernite, Max Ryabinin, Wietse de Vries, Hady Elsahar, Manan Dey, Shanya Sharma, Sampo Pyysalo, Veronika Laippala, Logesh Kumar Umapathi, Archiki Prasad, Elizabeth Keleshian, Shiyue Zhang, Fatma Tarlaci, Tal Perry, Ona de Gibert Bonet, Gunjan Chhablani, Madison May, Christophe Cerisara, Jonathan Chang, Arnaud Stiegler, Laurent Besacier, Nicolas Hervé, Tasnim Mohiuddin, Filip Ginter, Vinay Uday Prabhu, Trishala Neeraj, Matteo Manica, Tosin Adewumi, Rishi Bommasani, Fredrik Olsson, Sadid A. Hasan, Huu Nguyen, Gérard Dupont, Antoine Laurent, Salah Aït-Mokhtar, David I. Adelani, Dimitrios Christofidellis, Shamsuddeen Muhammad, Gabriel Ilharco, Alberto Bégue, Sebastian Nagel, Alessandro Suglia, Trieu Le, Emily Reif, Tolga Bolukbasi, Fabrice Zapfack, Thomas Wang, Ifeoluwatayo Ige, Max Huang, Kyle Lo, Dean Pleban, Idris Abdulmumin, Chien Vu, Aaron Gokaslan, Aashiq Muhamed, Leo Gao, Taewoon Kim, Laurel Orr, Sidd Karamcheti, Hieu Tran, Francesco De Toni, Long Phan, Binh Nguyen, Duong A. Nguyen, Edward Tan, Hieu Tran, Thanh Le, Trieu Le, Manuel Romero, Maria Grandury, Yulong Wang, Kenneth Heafield, Tong Zhou, Joppe Geluykens, María Grandury, Javier de la Rosa, Eduardo González Ponferrada, Vikas Raunak, Ebenge Usip, Aaron Gokaslan, Habiba Dahmani, Diwakar Mahajan, Chenghao Mou, Olivier Nguyen, Jean Louis Queguiner, Birger Moëll, Maraim Masoud, Zaid Alyafeai, Adela Ljajić, Suhas Pai, Levkovizh Efrat, Ariel Kreisberg Nitzav, Eyal Bar Natan, Stanislav Silberberg, Khalid Almubarak, Ian Yu, Leon Weber, Christopher Akiki, Jenny Chim, Hugo Laurençon, Jian Zhu, Rishi Bommasani, Jason Fries, Manan Dey, Natasha Seelam, Nora Kassner, Jörg Frohberg, Waris Radji, Daniel Whitenack, Shayne Longpre, Levkovizh Efrat, Sainyam Galhotra, Keith Stevens, Samuel Cahyawijaya, Jordan Clive, Alya Karim



Current chair(s)

Colin Raffel, Albert Villanova, Sampo Pyysalo, Filip Ginter

Current status

Scoping out further collaborative tasks within the group.

3. Modeling approach

Scope

The high level goal is to answer in natural language any prompted query expressed in natural language in a 0-shot manner:

- (a) Defining the prompted data creation and training procedure to make prompting less “accidental” than in GPT3
- (b) Looking at untapped sources of signal such as metadata and context and potential tools (such as retrieval) to include these signals
- (c) Adapting these in the context of multilinguality
- (d) Implementing and potentially small shared manual labor to create prompting instances
- (e) Coordinating with dataset and engineering working groups and training the beast
- (f) Coordinating with the evaluation working group to evaluate the trained model(s)

Resources - onboarding - documentation

Folder on the google drive: [Modeling documents](#) and in particular the [Entry point document](#)

Current members

Stephen Bach, M Saiful Bari, Stas Beckman, Iacer Calixto, Marine Carpuat, Joe Davison, Thibault Févry, Karën Fort, Jason Fries, Matthias Gallé, Claire Gardent, Braden Hancock, Yacine Jernite, Francois Lagunas, Veronika Laippala, Teven Le Scao, André Martins, Aurélie Névéol, Pierre-Yves Oudeyer, Matt Peters, Sampo Pyysalo, Colin Raffel, Nazneen Rajani, Verena Rieser, Sasha Rush, Victor Sanh, Timo Schick, Samson Tan, Eric Wallace, Thomas Wolf, Chen Xing, Caiming Xiong, Canwen Xu, Francois Yvon, Alexandre Allauzen, Yozh (jheuristic), Patrick Lewis, Peter Bloem, Josep Crego, Max Ryabinin Mohamed Ahmed, Manan Dey, Shanya Sharma, Iiro Rastas, Minh Quang Pham, Jean-Baptiste Cordonnier, Thijs Vogels, Ryan Teehan, Akash Mehra, Sandip Patil, Ruochen Zhang, Logesh Kumar Umapathi, Jonathan Berant, Xiang Zhou, Archiki Prasad, Fatma

Tarlaci, Antoine Bosselut, Edoardo M. Ponti, Nicola De Cao, Han Wang, Pasquale Minervini, Marius Mosbach, Nils Reimers, Eliza Szczechla, Germán Kruszewski, Patrick Fernandes, Albert Webson, Arnaud Stiegler, Tian Yun, Leandro von Werra, Gunjan Chhablani, Madison May, Christophe Cerisara, Iz Beltagy, Md Rabiul Awal, Pedro Henrique Martins, Olivier Ferret, Sylvain Viguier, Amar Viswanathan, Chenglei Si, Mike Tian-Jian Jiang, Kamalkumar Rathinasamy, Kaustubh Dhole, Jonathan Chang, Mohit Bansal, Luou Wen, Nora Kassner, James Briggs, Roman Castagné, Diwakar Mahajan, Thomas Wang, Yujian Gan, Pawan Sasanka Ammanamanchi, Jan-Christoph Kalo, Nicolas Chauville, Ben Peters, Abheesht Sharma, Tasnim Mohiuddin, Harshit Pandey, Urmish Thakker, Sushil Bharati, Rabin Banjade, Filip Ginter, Gorka Labaka, Elizabeth Keleshian, Siddharth Karamcheti, Vinay Uday Prabhu, Dimitra Gkatzia, Trishala Neeraj, Vishakha Sharma, Devamanyu Hazarika, Shubham Agarwal, Shannon Shen, Stella Biderman, Nickil Maveli, Sheng Shen, Jesse Dodge, Sadid A. Hasan, Ethan Perez, Stella Biderman, Taewoon Kim, Priyank Soni, Antoine Chaffin, Vikas Raunak, Zaid Alyafeai, Satviki Pathak, Debajyoti Datta, Zheng Xin Yong, Alessandro Suglia, Zhiqing Sun, Haryo Akbarianto Wibowo, Fabrice Zapfack, Lucile Saulnier, Josh Levy-Kramer, Anant Khandelwal, Ofir Press, Wilson Lee, Jake Tae, Arman Cohan, Trieu Le, Chien Vu, Tian Yun, Rose E. Wang, Aaron Gokaslan, Eric Hallahan, Leo Gao, Aashiq Muhamed, Suhas Pai, Yulong Wang, Alham Fikri Aji, Hieu Tran, Alyssa Loo, Alex Muzio, Long Phan, Duong A. Nguyen, Max Bartolo, Phil Wang, Manuel Romero, Olivia(Simin) Fan, Birger Moëll, Adela Ljajić, Eduardo González Ponferrada, Leon Weber, Khalid Almubarak, Thomas Metcalfe, Xiangru Tang, Isaac Rodríguez, Vlad Lialin, Shayne Longpre, Jordan Clive, Huy Ngo, Christopher Akiki, Newton Cheng, Varsha Jain

Current Area-chair(s)

Victor Sanh & Stephen Bach

Current status

Initial experiments (small scale quantitative experiments to compare different options for injecting promoting behaviors) + kick started the sub working groups from the kickoff brainstorming

Sub working groups

We organized the modeling approach working group into task forces (or sub working groups) around specific questions:

- [Architecture design](#)

Scope: Comparing and deciding which Transformer architecture would be best. Axis of choices include encoder-decoder/decoder-only/prefix-LM, long range attentions, etc.

- [Metadata](#)
Scope: Metadata gives information about the situation in which a text was written and is a source of signal we want to squeeze out to condition text on available metadata as an additional source of context.
- [Multilinguality](#)
Scope: We are interested in ensuring the model still has non-trivial few/O-shot abilities beyond English.
- [Prompt engineering](#)
Scope: We are interested in the emergence of 0-shot prompting behaviors in very large language models and in making 0-shot prompting abilities more robust.
- [Retrieve](#)
Scope: Inputs alone often don't contain enough information and providing additional context is required to faithfully react to the prompt without hallucinating or memorizing random things. We are interested in exploring whether a non-parametric component to the model would make sense.
- [Scaling-laws](#)
Scope: We are interested in estimating the best model, data and compute size configuration before scaling to a very large model so that we make the best use of the computational resources.

3.1. Modeling – Architecture and Scaling

Scope


Comparing and deciding which Transformer architecture would be best. Axis of choices include encoder-decoder/decoder-only/prefix-LM, long range attentions, etc. Also, estimating the best model, data and compute size configuration before scaling to a very large model so that we make the best use of the computational resources.

Resources - onboarding - documentation

Folder on the google drive: [Modeling documents](#)

Current members

Thomas Wolf, Victor Sanh, Angela Fan, François Lagunas, Teven Le Scao, Sasha Rush, Yacine Jernite, Colin Raffel, Pasquale Minervini, Matthias Gallé, François Yvon, Gunjan Chhablani, M Saiful Bari, Timo Schick, André Martins, Suzana Ilic, Stephen Bach, Ryan Teehan, Arnaud Stiegler, Germán Kruszewski, Thomas Wang, Eliza Szczechla, Urmish Thakker, Sushil Bharati, Vinay Uday Prabhu, Vishakha Sharma, Nimshi Venkat Meripo, Ruochen Zhang, Shubham Agarwal, Shannon



Shen, Nickil Maveli, Sheng Shen, Matthew Peters, Jesse Dodge, Stella Biderman, Jonathan Chang, Marius Mosbach, Jake Tae, Pawan Sasanka Ammanamanchi, Priyank Soni, Eric Wallace, Matt Peters, Jonathan Berant, Lewis Tunstall, Siddharth Karamcheti, Tasnim Mohiuddin, Trishala Neeraj, Ahmed Baruwa, Jason Phang, Gabriel Ilharco, Canwen Xu, Taewoon Kim, Srulik Ben David, Iacopo Poli, Alberto Bégué, Dimitri Lozeve, Alessandro Suglia, Zhiqing Sun, Sandip Patil, Zheng Xin Yong, Huu Nguyen, Anant Khandelwal, Ofir Press, Arman Cohan, Christopher Grainger (he/him), Manan Dey, Roman Castagne, Seungju Han, Eric Hallahan, Leo Gao, Conglong Li, Jin Yong Yoo, Aashiq Muhamed, Corby Rosset, Trent Oh, Anna Shvets, Suhas Pai, Alham Fikri Aji, Diwakar Mahajan, Alexandre Muzio, Simeng Han(simeng.han@yale.edu), Chunting Zhou, Max Bartolo, Gurminder Singh, Manuel Romero, Birger Moëll, Adela Ljajić, Eduardo González Ponferrada, Thomas Metcalfe, Dragomir Radev, Xiangru Tang, Isaac Rodríguez, Abhilash Majumder, Jordan Clive, Huy Ngo, Christopher Akiki, Tri Dao, Khalid Almubarak, Niklas Muennighoff, Axel Marmet, Antonio Domínguez, Varsha Jain, Rishabh Anand

Current chair(s)

Julien Launay, Iz Beltagy

Current status

Merged Architecture Design and Scaling Laws into one group, Architecture and Scaling.

3.2. Modeling – Metadata

Scope

Metadata gives information about the situation in which a text was written and is a source of signal we want to squeeze out to condition text on available metadata as an additional source of context.

Resources - onboarding - documentation

Folder on the google drive: [Modeling documents](#)

Current members

Thomas Wolf, Victor Sanh, Angela Fan, Teven Le Scao, Yacine Jernite, Colin Raffel, Patrick Lewis, Manan Dey, Matthias Gallé, M Saiful Bari, Timo Schick, Suzana Ilic, Stephen Bach, Eliza Szczechla, Kaustubh Dhole, Urmish Thakker, Vishakha Sharma, Shannon Shen, Gérard Dupont, Jonathan Chang, Marius Mosbach, Huu Nguyen, Taewoon Kim, Shanya Sharma, Alessandro Suglia, Mike Tian-Jian Jiang, Lucile Saulnier, Christopher Klammer, Anant Khandelwal, Chien Vu, Mounir Boudali,

Charles Welch, Jan-Christoph Kalo, Raghav Tandon, Jason Phang, Aaron Gokaslan, Adela Ljajić, Anna Wegmann, Jordan Clive, Huy Ngo, Antonio Domínguez

Current chair(s)

Timo Schick, Nora Kassner

Current status

Recently created

3.3. Modeling – Multilinguality

Scope

We are interested in ensuring the model still has non-trivial few/0-shot abilities beyond English.

Resources - onboarding - documentation

Folder on the google drive: [Modeling documents](#)

Current members

Thomas Wolf, Victor Sanh, Angela Fan, Teven Le Scao, Sasha Rush, Yacine Jernite, Colin Raffel, Patrick Lewis, Jonas Pfeiffer, Manan Dey, Matthias Gallé, M Saiful Bari, Nils Reimers, Suzana Ilic, Stephen Bach, Edoardo Ponti, Mohamed Ahmed, Archiki Prasad, Shanya Sharma, Eliza Szczechla, Patrick Fernandes, Tasnim Mohiuddin, Kaustubh Dhole, Urmish Thakker, David Stap, Vinay Uday Prabhu, Vishakha Sharma, Lintang Sutawika, Md Rabiul Awal, Sadid A. Hasan, Iacer Calixto, Ahmed Baruwa, Jonathan Chang, Samson Tan, Bjarke Felbo, Vikas Raunak, M Ali Jauhar, Zheng Xin Yong, Evangelia Gogoulou, Sandip Patil, Pedro Ortiz Suárez, Anant Khandelwal, Trieu Le, Evgeniia Razumovskaia, Thomas Wang, Alexandre Berard, Chien Vu, Leo Gao, Taewoon Kim, Diwakar Mahajan, Alham Fikri Aji, Alyssa Loo, Alexandre Muzio, Max Bartolo, Habiba Dahmani, Manuel Romero, Birger Moëll, Leon Weber, Khalid Almubarak, Shachar Mirkin, Abhilash Majumder, Davut Emre Taşar, Shayne Longpre, Aashiq Muhamed, Huy Ngo, Newton Cheng, Hailey Schoelkopf, Jungo Kasai, Dragomir Radev, Ali Hürriyetoğlu, Genta Indra Winata, Etsuko Ishii

Current chair(s)

Vassilina Nikoulina, Hady Elsahar

Current status

Recently created

3.4. Modeling – Prompt engineering

Scope

We are interested in the emergence of 0-shot prompting behaviors in very large language models and in making 0-shot prompting abilities more robust.

Resources - onboarding - documentation

Folder on the google drive: [Modeling documents](#)

Current members

Thomas Wolf, Victor Sanh, Angela Fan, François Lagunas, Teven Le Scao, Joe Davison, Sasha Rush, Yacine Jernite, Colin Raffel, Pasquale Minervini, Canwen Xu, Iz Beltagy, Manan Dey, Matthias Gallé, Claire Gardent, Pierre-Yves Oudeyer, Eric Wallace, Gunjan Chhablani, M Saiful Bari, Jonathan Berant, Y-Lan Boureau, Braden Hancock, Suzana Ilic, Stephen Bach, Jason Fries, Ruiqi Zhong, Mohamed Ahmed, Peter Bloem, Arnaud Stiegler, Marius Mosbach, Germán Kruszewski, Thomas Wang, Shanya Sharma, Eliza Szczechla, Madison May, Nicolas Chauville, Kaustubh Dhole, Urmish Thakker, Mike Tian-Jian Jiang, Dimitra Gkatzia, Tasnim Mohiuddin, Vishakha Sharma, Ruochen Zhang, Devamanyu Hazarika, Abheesht Sharma, Shannon Shen, Sheng Shen, Michael McKenna, Harshit Pandey, Ethan Perez, Jake Tae, Samira Shaikh, Taewoon Kim, Albert Webson, Priyank Soni, Satviki Pathak, Arun Raja, Debajyoti Datta, Trishala Neeraj, Vikas Raunak, Zheng Xin Yong, Patrick Lewis, Rachel Bawden, Andrea Santilli, Lintang Sutawika, Ryan Teehan, Matteo Manica, Animesh Prasad, Han Wang, Nihal Nayak, Zaid Alyafeai, Fatma Tarlaci, Michael McKenna, Srulik Ben David, Jacopo Staiano, Jonathan Chang, Zhiqing Sun, Kaustubh Dhole, Antoine Chaffin, Jos Rozen, Sandip Patil, Bec Johnson, Max Huang, Anant Khandelwal, Tian Yun, Wojciech Kusa, Vikas Raunak, Julien Launay, Tali Bers, Thibault Fevry, Nafis Abrar, Seungju Han, Stella Biderman, Eric Hallahan, Leo Gao, Jin Yong Yoo, Aashiq Muhamed, Jason Phang, Adam Roberts, Diwakar Mahajan, Alham Fikri Aji, Natasha Seelam, Edward Tan, Ebenge Usip, Simeng Han, Joppe Geluykens, Max Bartolo, Gurvinder Singh, Shaked Brody, Olivia(Simin) Fan, Lina Rojas-Barahona, Birger Moëll, Mihir Kale, Dragomir Radev, Leon Weber, Khalid Almubarak, Xiangru Tang, Vlad Lialin, Reshith Adithyan, Niall Taylor, Haobo Yuan, Tianyang Li, Yi Zhang, Andrey Kormilitzin, Jordan Clive, Huy Ngo, Tom Liu, Rajaswa Patil, Newton Cheng, Hailey Schoelkopf, Tomáš Mikula, Antonio Domínguez, Varsha Jain



Current chair(s)

No chair yet

Current status

Recently created

3.5. Modeling – Retrieve

Scope

Inputs alone often don't contain enough information and providing additional context is required to faithfully react to the prompt without hallucinating or memorizing random things. We are interested in exploring whether a non-parametric component to the model would make sense.

Resources - onboarding - documentation

Folder on the google drive: [Modeling documents](#)

Current members

Thomas Wolf, Victor Sanh, Lysandre Debut, Angela Fan, Teven Le Scao, Sasha Rush, Yacine Jernite, Colin Raffel, Patrick Lewis, Matthias Gallé, Minh Quang Pham, M Saiful Bari, Timo Schick, Nils Reimers, Suzana Ilic, Stephen Bach, Nicola De Cao, Marius Mosbach, Thomas Wang, Eliza Szczechla, Madison May, Jan-Christoph Kalo, Urmish Thakker, Siddharth Karamcheti, Tasnim Mohiuddin, Vishakha Sharma, Ruochen Zhang, Devamanyu Hazarika, Nickil Maveli, Iacer Calixto, Jonathan Chang, Taewoon Kim, Vikas Raunak, Animesh Prasad, Dimitri Lozeve, Jakub Zavrel, Zhiqing Sun, Manan Dey, Sandip Patil, Antoine Chaffin, Max Huang, Anant Khandelwal, Bo Wang, Wojciech Kusa, Arman Cohan, Ebenge Usip, Trent Oh, Tanmay Laud, Seungju Han, Jin Yong Yoo, Aashiq Muhamed, Jason Phang, Suhas Pai, Zheng-Xin Yong, Shanya Sharma, Junxian He, Chunting Zhou, Gurminder Singh, Olivia(Simin) Fan, Birger Moëll, Leon Weber, Thomas Metcalfe, Jordan Clive, Clara Bersch, Max Tian, Huy Ngo, Newton Cheng, Niklas Muennighoff, Axel Marmet, Sameera Horawalavithana

Current chair(s)

Nils Reimers



Current status

Recently created

4. Tokenization

Scope

Define and develop the tokenization used by the model:

- the tokenization algorithm itself
- the tokenization in the context of multilinguality
- implement it and coordinate with the dataset and modeling working groups to prepare the data for training

Resources - onboarding - documentation

Not resource at the moment

Current members

M Saiful Bari, Anthony Moi, Victor Sanh, Samson Tan, Francois Yvon, Wietse de Vries, Mohamed Ahmed, Wilson Lee, Sandip Patil, Rico Sennrich, Shiyue Zhang, Maxime De Bruyn, Vassilina Nikoulina, Zaid Alyafeai, Ben Peters, Chenglei Si, Mike Tian-Jian Jiang, Kamalkumar Rathinasamy, Kaustubh Dhole, Benoît Sagot, Roman Castagné, Thomas Wang, Uğur Ünal, Md Rabiul Awal, Tasnim Mohiuddin, Rachel Bawden, Elizabeth Keleshian, Manan Dey, Shannon Shen, Sadid A. Hasan, Arun Raja, Priyank Soni, Bjarke Felbo, Sabrina Mielke, Shamsuddeen Muhammad, Taewoon Kim, Zhiqing Sun, Elizabeth Salesky, Fabrice Zapfack, Ifeoluwatayo Ige, Anant Khandelwal, Shachar Mirkin, Hugo Sonnery, Benjamin Heinzerling, Christopher Grainger (he/him), Phoebe Mulcaire, Deepak Tunuguntla, Phil Wang, Birger Moëll, Adela Ljajić, Jonathan Chang, Yuval Pinter, Davut Emre Taşar, Aashiq Muhamed, Christopher Akiki, Samuel Cahyawijaya, Antonio Domínguez, Rishab Anand

Current chair(s)

Benoît Sagot, Samson Tan

Current status

Kicked off the discussion

5. Evaluation

Scope

This working group focuses on making the question precise: What makes a language model “good”? We aim to:


- Develop standardized measures for evaluating models both 1) relative to some extrinsic gold standard (e.g., human performance/theoretical bounds) and 2) relative to one another
- Design a diverse, multifaceted suite of evaluations that represents 1) the different ways in which people will use language models in practice and 2) the different scientific communities who are interested in the capabilities of language models

Resources - onboarding - documentation

Our onboarding document is [here](#).

Current members

Laurent Besacier, Iacer Calixto, Marine Carpuat, Joe Davison, Karën Fort, Yacine Jernite, Teven Le Scao, Aurélie Névél, Verena Rieser, Sasha Rush, Victor Sanh, Thomas Wolf, Francois Yvon, Benoît Crabbé, Pierre Colombo, Marie Candito, Yozh (jheuristic), Wietse de Vries, Maite Melero, Antoine Neuraz, Hady Elsahar, Vassilina Nikoulina, Manan Dey, Shanya Sharma, Philippe Muller, Wilson Lee, Ryan Teehan, Akash Mehra, Eric Wallace, Xiang Zhou, Saad Mahamood, Aline Villavicencio, Archiki Prasad, Jan-Christoph Kalo, Nora Kassner, Shiyue Zhang, Olga Majewska, Chrysoula (Chryssa) Zerva, Taisiya (Taya) Glushkova, Jos Rozen, Maraim Masoud, Yonatan Belinkov, Pierre-Yves Oudeyer, Marius Mosbach, Miruna Clinciu, Albert Webson, Tian Yun, Leandro von Werra, Salah Aït-Mokhtar, Gunjan Chhablani, Md Rabiul Awal, Olivier Ferret, Amar Viswanathan, Mike Tian-Jian Jiang, Kaustubh Dhole, Mohit Bansal, Abhinav Ramesh Kashyap, Caroline Brun, Luou Wen, Debajyoti Datta, Namrata Mukhija, Diwakar Mahajan, Samson Tan, Thomas Wang, Julien Tourille, Andrea Santilli, Nicolas Chauville, Pawan Sasanka Ammanamanchi, Abheesht Sharma, José Hernández-Orallo, Harshit Pandey, Urmish Thakker, Rabin Banjade, Rachel Bawden, Gorka Labaka, Dimitra Gkatzia, Elizabeth Keleshian, Trishala Neeraj, Ruochen Zhang, Alison Callahan, Shubham Agarwal, Jaap Jumelet, Arun Raja, Vered Shwartz, Nickil Maveli, Veronika Laippala, Sadid A. Hasan, Oskar van der Wal, Dan Garrette, Ethan Perez, Rodrigo Wilkens, Ahmed Baruwa, Satviki Pathak, Vikas Raunak, Vincent Claveau, Nafise Sadat Moosavi, Shamsuddeen muhammad, Giada Pistilli, Bec Johnson, Zeerak Waseem, Jekaterina Novikova, Shachar Mirkin, Kerenza Doxolodeo, Rose E. Wang, Karim Lasri, Jonathan Bragg, Leo Gao, Eric Hallahan, Deepak



Tunuguntla, Jason Phang, Corby Rosset, Taewoon Kim, Suhas Pai, Weijia Xu, Max Bartolo, Zsombor Koman, Birger Moëll, Haoming Jiang, Thomas Metcalfe, Omer Antverg, Shira Guskin, Jordan Clive, Aashiq Muhamed, Huy Ngo, Christopher Akiki, Newton Cheng, Hailey Schoelkopf, Jungo Kasai, Matthias Samwald, Genta Indra Winata, Tomáš Mikula, Antonio Domínguez

Current chair(s)

Aurélie Névéol, Dan Garrette, Ellie Pavlick, Marine Carpuat, Thomas Scialom, Sebastian Gehrmann, Verena Rieser

Current status

Our subgroups have continued to refine their focus (see updates below). As a full group, we are working on populating a spreadsheet with an extensive (if not comprehensive) list of existing evaluation resources along with key metadata on the research questions and design decisions underlying each. The goal is to release the resulting spreadsheet to the community as a useful documentary resource, and then use it ourselves as the basis for our baseline evaluation suite. We welcome contributions from the community to the spreadsheet: the dataset tracker is [here](#) and the form for submitting new datasets is [here](#).

Our goal is to filter this list down to a small, interpretable set of evaluations that will make up our baseline evaluation suite. We plan to finalize the spreadsheet in the coming week or two, and then spend several weeks finalizing a short list of evaluations. Our intent is for the short list to be “quality over quantity”—we want to include a smaller number of datasets such that changes in performance are easily interpretable, rather than a large number of possibly redundant or highly-correlated metrics of performance.

5.1. Evaluation – Extrinsic Evaluation

Scope

Standardize a suite of application-oriented downstream tasks (e.g., Question Answering, Summarization, Machine Translation) on which to evaluate models

Resources - onboarding - documentation

[Google drive sub-folder for evaluation](#)

[Main on-boarding document.](#)

Current members

Shiyue Zhang (she/her), Pawan Sasanka Ammanamanchi (he), Salah Aït-Mokhtar, Julien Tourille, Luou Wen (any), Rachel Bawden, Joe Davison (he/him), Radu Florian, Thomas Scialom, Verena Rieser, Diwakar Mahajan, Sadid A. Hasan, Vered Shwartz, Rabin Banjade, Yacine Jernite, Iacer Calixto, Veronika Laippala, MaraimM, Md Rabiul Awal, Yonatan Belinkov, Samira Shaikh, Vikas Raunak, Vincent Claveau, M Saiful Bari, Miruna Clinciu, Nafise Sadat Moosavi, Ahmed Baruwa, Ksenia Shkaruta, Manan Dey, Gabriele Sarti, Karim Lasri, Jason Phang, Taewoon Kim, Suhas Pai, Leon Weber, Shayne Longpre, Jordan Clive, Aashiq Muhamed, Huy Ngo, Christopher Akiki, Newton Cheng, Yada Pruksachatkun

Current chair(s)

Thomas Scialom, Verena Rieser, Ehud Reiter

Current status

Recently created

5.2. Evaluation – Intrinsic Evaluation

Scope

Standardize a suite of evaluations focused on interpretability of internal linguistic representations, including but not limited to linguistic knowledge (syntax, morphology, semantics), commonsense/world knowledge, and visualizations

Resources - onboarding - documentation

[Google drive sub-folder for evaluation](#)

[Main on-boarding document.](#)

Current members

Miruna-Adriana Clinciu, Amar Viswanathan, Reshinh, Nora Kassner (she/her), Inyoung, Gunjan Chhablani (he/him), Harshit Pandey (He/Him), Abheesht Sharma (he/him), Olga Majewska (she/her), Aline Villavicencio (she/her), Jan-Christoph Kalo (he/him), Shubham Agarwal, Jaap Jumelet (he/his), Eleni Metheniti (she/her), Saad Mahamood, Albert Webson, Sebastian Gehrmann, Rishi Bommasani, Sadid A. Hasan, Yanai Elazar, Samira Shaikh, Rodrigo Wilkens, Satviki Pathak, M Saiful Bari, Yonatan Belinkov, Gabriele Sarti, Ksenia Shkaruta, Blanca Calvo

Figueras, Manan Dey, Jelle Zuidema, Rachit Bansal, Karim Lasri, Taewoon Kim, Simeng Han, Habiba Dahmani, Thomas Metcalfe, Yuval Pinter, Jordan Clive, Aashiq Muhamed, Joseph Imperial, Huy Ngo, Christopher Akiki, Rajaswa Patil, Newton Cheng

Current chair(s)

Sebastian Gehrmann

Current status

Recently created

5.3. Evaluation – Few-Shot Generalization

Scope

Standardize a set of tasks and procedures for evaluating models' ability to generalize outside of the training domain/distribution

Resources - onboarding - documentation

[Google drive sub-folder for evaluation](#)

[Main on-boarding document.](#)

Current members

Arun Raja, Trishala Neeraj, Albert Webson (he/they), Md Rabiul Awal (he/him), Pierre Colombo, Urmish Thakker, Mike Tian-Jian Jiang (he/him), Abhinav Ramesh kashyap, Xiang Zhou (he/him/his), Dimitra GKATZIA, Wilson Lee (he/him), Archiki Prasad (she/her), Andrea Santilli (he/him), Debajyoti Datta, Jose Hernandez-Orallo, Ryan Teehan (he/him), Jos Rozen (he-him), Namrata Mukhija (she/her), Tian Yun (He/Him), Ellie Pavlick, Ruochen Zhang (she/her), Sadid A. Hasan, Yanai Elazar, Jekaterina Novikova, Vincent Claveau, M Saiful Bari, Gonzalo Jaimovitch-López (He/Him,She/Her,They/Them), Aakanksha Naik, Taewoon Kim, Alessandro Suglia, Max Huang, Arman Cohan, Jake Tae, Jonathan Bragg, Jin Yong Yoo, Jason Phang, Adam Pearce, Haoming Jiang, Shira Guskin, Jordan Clive, Aashiq Muhamed, Huy Ngo, Newton Cheng, Francesco De Toni

Current chair(s)

Ellie Pavlick



Current status

We are working on populating the shared Evaluation WG spreadsheet with datasets that are aimed at evaluating few-shot learning and generalization. Our goal is to begin filtering these datasets to a short list in the following week. We expect that some work will be required to adapt datasets to a few-shot setting (e.g., combining a Wikipedia training set with a biomedical domain test set to test few-shot domain transfer).

5.4. Evaluation – Bias, Fairness, and Social Impact

Scope

Standardize a set of tasks and procedures for evaluating models' biases and risks, with focus on the impact model will have on people if deployed in real-world scenarios.

Resources - onboarding - documentation

[Google drive sub-folder for evaluation](#)

[Main on-boarding document.](#)

Current members

Ioana Baldini, Hady Elsahar (he/him), Rabin Banjade, Verena Rieser (she/her), Samson Tan (he/him), Brun Caroline, Aurelie Neveol, Diwakar Mahajan, Rishi Bommasani, Oskar van der Wal, Deepa Muralidhar, Maxime Amblard, Jake Tae (he/him), Vincent Claveau, M Saiful Bari, M Ali Jauhar, Marie Candito, Miles Brundage, Debajyoti Datta, Maraim Masoud, Xudong Shen, Miruna Clinciu, Giada Pistilli, Bec Johnson, Arjun Subramonian (they/them), Zeerak Waseem, Valentin Danchev, Manan Dey, Shanya Sharma, Trishala Neeraj, Mounir Boudali, Jason Phang, Anaelia Ovalle, Taewoon Kim, Deepak Tunuguntla, Himanshu Beniwal, Dongyeop Kang, Isaac Rodríguez, Shayne Longpre, Aashiq Muhamed, Newton Cheng, Sarah Wyer, Rodrigo Beceiro, Hailey Schoelkopf, Yada Pruksachatkun, Jordan Clive

Current chair(s)

Aurelie Neveol

Current status

Recently created

5.5. Evaluation – Multilinguality

Scope

Interface with all other subgroups to ensure all proposed evaluations cover a sufficient range of languages and language families, and potentially design other evaluations specific to multilinguality

Resources - onboarding - documentation

[Google drive sub-folder for evaluation](#)

[Main on-boarding document](#)

Current members

Maraim Masoud, Laurent Besacier, Yacine Jernite, Shanya (She/Her), Wietse de Vries (he/him/his), Diwakar mahajan(he/him), Manan Dey (He/Him), Marine Carpuat, Iacer Calixto, Fatma Tarlaci, Lintang Sutawika (He/Him), Evi Gogoulou (she/her), Aline Villavicencio, David I. Adelani, M Saiful Bari, Rachel Bawden, Samson Tan, Md Rabiul Awal, Vassilina Nikoulina, Rabin Banjade, M Ali Jauhar, Ahmed Baruwa, Rui Zhang, Trieu Le, Omar Sanseviero, Alham Fikri Aji, E. Razumovskaia, Jekaterina Novikova, Rua Ismail, Tiago Torrent, Wilson Wongso, Yong (Zheng Xin Yong), Zaid Alyafeai, Idris Abdulmumin, Kerenza Doxolodeo, Daniil Larionov, Hady Elsahar, Victoria Lin, Deepak Tunuguntla, Taewoon Kim, Dragomir Radev, Joydeep Bhattacharjee, Johannes Otterbach, Shayne Longpre, Khalid Almubarak, Aashiq Muhamed, Huy Ngo, Newton Cheng, Ifeoluwatayo A. Ige, Hailey Schoelkopf, Jungo Kasai, Ali Hürriyetoğlu, Genta Indra Winata

Current chair(s)

Marine Carpuat, Dan Garrette

Current status

Recently created

6. Interpretability and interaction/visualization

Scope

Developing tools and interactive methods horizontally to the other working groups to:

- help debugging model (see <https://exbert.net>)
- help identifying bias and mitigate for it
- create interactive prototypes for human-AI generation to test model under human constraints
- encourage to train models with more interpretable latents
- more...

Resources - onboarding - documentation

No resources at the moment

Current members

Hendrik Strobelt, Terrell Ibanez, Pierre Colombo, Mohamed Ahmed, Ryan Teehan, Ioana Baldini, Yonatan Belinkov, Miruna Clinciu, Eliza Szczechla, Patrick Fernandes, Gunjan Chhablani, Mohit Bansal, Debajyoti Datta, Sebastian Gehrmann, Diwakar Mahajan, Luca Moschella, Andrea Santilli, Uğur Ünal, Abheesht Sharma, Urmish Thakker, Sushil Bharati, Trishala Neeraj, Jaap Jumelet, Shannon Shen, Nickil Maveli, Michael McKenna, Matteo Manica, Sara Meftah, Sadid A. Hasan, Taewoon Kim, Satviki Pathak, James Wexler, Rishi Bommasani, Nora Kassner, Manan Dey, Germán Kruszewski, Gabriele Sarti, Vinay Prabhu, Maraim Masoud, Emily Reif, Luciana Benotti, Tolga Bolukbasi, Max Huang, Shachar Mirkin, Siddharth Karamcheti, Natasha Seelam, Trent Oh, Vadim Fomin, Oleg Serikov, Antoine Bosselut, Rose E. Wang, Stella Biderman, Rachit Bansal, Karim Lasri, Jin Yong Yoo, Aaron Gokaslan, Zsombor Koman, Christopher Akiki, Adam Pearce, Isaac Rodríguez, Rajaswa Patil, Samuel Cahyawijaya, Ruochen Zhang, Khalid Almubarak, Nick Doiron, Newton Cheng, Sarah Wyrer, Ethan Kim

Current chair(s)

Hendrik Strobelt, Ryan Teehan

Current status

Kicked off the discussion

7. Engineering/scaling

Scope

A working group focusing on the technical challenges of training at scale on several hundred GPUs and how to make the best use of the (large) compute budget we have.

Note that participating in this working group does not imply that you will have direct access to the supercomputer since there are additional (quite strong) national restrictions on the access to this machine (see some details in the [section on access to compute here](#)). It does mean however that you will participate in the discussion on these aspects.

Resources - onboarding - documentation

Current members

Sean Narenthiran, Stas Bekman, Thomas Wolf, Teven Le Scao, Morgan Funtowicz, Yozh (justheuristic), Max Ryabinin, Jean-Baptiste Cordonnier, Thijs Vogels, Akash Mehra, Antoine Bosselut, Samyam Rajbhandari, Gérard Dupont, Minjia Zhang, guillaume Alleon, Christophe Cerisara, James Briggs, Uğur Ünal, Harshit Pandey, Urmish Thakker, Sushil Bharati, Myriam Peyrounette, Hatim Bourfoune, Hadrien de Dompure, Kamel Guerda, Siddharth Karamcheti, Laurel Orr, Espoir Murhabazi, Thomas Wang, Luca Moschella, Francois Courteille, Kari Briski, Meriem Bendris, Mohammad Shoeybi, Bryan Catanzaro, Jared Casper, Joey Conway, Sylvain Jauegy, Camille Parisel, Deepak Narayanan, Bertrand Cabot, Rémi Lacroix, Sampo Pyysalo, Pierre-François Lavallée, Jeff Rasley, Olatunji Ruwase, Pawan Sasanka Ammanamanchi, Vikas Raunak, M Ali Jauhar, Arun Raja, Alberto Bégué, Dimitri Lozeve, Michael Z Wang, Zenodia Charpy, Shaden Smith, Josh Levy-Kramer, Rodrigo Beceiro, Aaron Gokaslan, Joseph Tobing, Oleg Serikov, Harm de Vries, Conglong Li, Seungju Han, Corby Rosset, Jason Phang, Eric Hallahan, Taewoon Kim, Suhas Pai, Yulong Wang, Jonathan Chang, Leo Gao, Alexandre Muzio, Gurvinder Singh, Tanmoy Sarkar, Dragomir Radev, Xiangru Tang, Raymond Li, Jordan Clive, Aashiq Muhamed, Hailey Schoelkopf, Muhtasham Oblokulov, Antonio Domínguez

Current chair(s)

Yozh (jheuristic), Max Ryabinin, Minjia Zhang



Current status

8. Collaborations with education and research communities

Scope

Around the workshop there are many actors with which we could do actions or events for collaborating or raising awareness.

Examples:

- Collaborating with ELLIS for student workshop
- Finding possibilities for early students (masters/licences) to be either involved in some part of the project or work on related projects (e.g. Hackathon)

Also it seems important to make sure we don't do some unreflected choices which can impact wide access to the workshop for the international community (typical example: selecting a location with VISA restrictions for an in-person event) (could be separated Working Groups).

Resources - onboarding - documentation

Current members

Thomas Wolf, Dieuwke Hupkes, Yozh (jheuristic), Omar Sanseviero, Peter Bloem, Manan Dey,, Andrea Santilli, Maraim Masoud, David Stap, Madhulika Srikumar, Maxime Amblard, Salomey Osei, Ksenia Shkaruta, Yonatan Belinkov, Christopher Klamm, Taewoon Kim, Tanmoy Sarkar, Dragomir Radev, Peter Troxler

Current chair(s)


Dieuwke Hupkes, Yozh (jheuristic), François Yvon

Current status

9. Carbon footprint

Scope

A Working Group interested in evaluating the carbon footprint of the collaborative tasks (model training, inference, evaluation, dataset creation, sharing) and proposing ways to reduce this footprint.



We will be working on this from two angles:

1. We will develop a plugin for transformers to automatically collect data about the energy consumption and carbon emission.
2. We will work together with modeling WG to provide high-quality compressed models (distillation, quantization, pruning etc.) thus can reduce the carbon footprint for downstream practitioners.

Resources - onboarding - documentation

Current members

Thomas Wolf, Anne-Laure Ligozat, Vassilina Nikoulina, Ari Jankelowitz, Elizabeth Keleshian, Han Wang, Sylvain Viguiet, Julien Launay, Uğur Ünal, Urmish Thakker, Vishakha Sharma, Benoît Sagot, Michael McKenna, Sheng Shen, Sasha Luccioni, Jesse Dodge, Jake Tae, Anthony Hevia, Wangchunshu Zhou, Manan Dey, Xabier Soto, Stella Biderman, Pawan Sasanka Ammanamanchi, Karim Lasri, Joseph Tobing, Taewoon Kim, Anthony Hevia, Peter Henderson, Johannes Otterbach, Samuel Cahyawijaya, Jordan Clive, Genta Indra Winata, Alya Karim, Maurice Jakesch

Current chair(s)

Anne-Laure Ligozat, Canwen Xu, Sasha Luccioni, Sylvain Viguiet

Current status

10. [Meta-Group] Social Impact Across Groups

Scope

While each of the working groups is responsible for considering the prospective social impacts and for researching the relevant social context for its decisions, these aspects of the various tasks cannot be considered in isolation from each other. The purpose of this working group will be to collaborate with each of the others to share skills and knowledge, propose general guidelines, identify common threads and relevant differences across their respective societal questions, and aggregate findings in the context of the overall effort.

This group is intended mostly as a **meta-group** where questions can be raised and discussed.



Resources - onboarding - documentation

Current members

Yacine Jernite, Thomas Wolf, Samson Tan, Omar Sanseviero, Peter Bloem, Karen Fort, Aurélie Névéol, Pierre-Yves Oudeyer, Christophe Cerisara, Benoît Sagot, Madhulika Srikumar, Sheng Shen, Rishi Bommasani, Aviv Ovadya, Huu Nguyen, Sasha Luccioni, Meheresh Masanpally, Shanya Sharma, Zeerak Waseem, Taewoon Kim, Pedro Ortiz Suárez, Debajyoti Datta, Maraim Masoud, Xudong Shen, Arjun Subramonian (they/them), Fabrice Zapfack, Christopher Klamm, Valentin Danchev, Sarah Wyer, Maurice Jakesch, Peter Troxler

Current chair(s)

Karen Fort, Michael McKenna, Margot Mieskes

Current status

11. Sharing and accessibility of model and dataset

Scope

Prepare policies and ethical framework for sharing the model; prepare tools and strategies to make the model easy to use and accessible to the community.

Resources - onboarding - documentation

Current members

Thomas Wolf, Alexandre Allauzen, Benoit Crabbé, Rajaram Sivaramakrishnan, Reshinh Adithyan, Benoît Sagot, Vishakha Sharma, Canwen Xu, Max Huang, Maxime Amblard, Shannon Shen, M Ali Jauhar, Miles Brundage, Christopher Klamm, Natasha Seelam, Valentin Danchev, Carlos Munoz Ferrandis, Rose E. Wang, David Lanksy, Aaron Gokaslan, Danish Contractor, Anima Shukla, Somaieh Nikpoor, Sergi Gálvez Duran, Giada Pistilli, Taewoon Kim, Harrison Pim, Anna Rogers, Sainyam Galhotra, Sarah Wyer

Current chair(s)

Thomas Wolf, Valentin Danchev, Danish Contractor and Aaron Gokaslan



Current status

12. Challenges in the biomedical domain

Scope

This transversal WG will focus on questions specifically related to the biomedical domain. It would include (non exhaustive list): data sourcing, modeling, evaluation, interpretability, data privacy,...

Ideally, it would include people from other WGs.

Resources - onboarding - documentation

Current members

Antoine Neuraz, Thomas Wolf, Aurélie Névéol, Vassilina Nikoulina, Salah Aït-Mokhtar, Jason Fries, Olivier Ferret, Diwakar Mahajan, Sushil Bharati, Laurel Orr, Byung-Hak Kim, Vishakha Sharma, Manan Dey, Alison Callahan, Debajyoti Datta, Maya Varma, Michael McKenna, Sadid A. Hasan, Iacer Calixto, Maxime Amblard, Vincent Claveau, Andrey Kormilitzin, Matteo Manica, Aakanksha Naik, Xabier Soto, Dimitri Lozeve, Bo Wang, Natasha Seelam, Wojciech Kusa, Kamal Raj, Tanmay Laud, Valentin Danchev, Taewoon Kim, Leon Weber, Zheng-Xin Yong, Nihal Nayak, Anaelia Ovalle, Xiangru Tang, Gurvinder Singh, Eduardo Gonzalez Ponferrada, Christopher McMaster, Ruisi Su, Gabriel Altay, Julio Bonis Sans, Matthias Samwald, Samuele Garda, Sunny Myungsun Kang, Milad Moradi, Simon Ott, Samuel Cahyawijaya, Ayush Singh, Trishala Neeraj, Uzay Macar, Alham Fikri Aji, Karthik Rangasai, Jordan Clive, Robert Martin, Shubhanshu Mishra, Michal Malyska, Francesco de Toni, Sameera Horawalavithana, Debora Marks, Florian Borchert

Current chair(s)

Jason Fries, Debajyoti Datta, Natasha Seelam

Current status

13. Language Modeling for Mathematics and Theorem Proving

Scope

Language models and other neural methods are increasingly used as parts of AI systems that automatically prove theorems and perform related tasks in informal and formal mathematics. The proposed working group would create and share large datasets, models and software tools in this AI field.

Initial goals are collecting a large public corpus of theorem proving and related data, as well as training a model that can serve as a publicly available pre-trained model, which is missing in the current situation.

Resources - onboarding - documentation

Not yet available, but coming in the future.

Entry document:

[Link](#) (placeholder document for now)

References:

These are some of the papers that are relevant to our efforts.

<https://easychair.org/publications/paper/g38n>

<https://ebooks.iospress.nl/publication/55039>

https://link.springer.com/chapter/10.1007%2F978-3-030-53518-6_24

https://link.springer.com/chapter/10.1007%2F978-3-319-94205-6_37

<https://arxiv.org/abs/2009.03393>

<https://arxiv.org/abs/2102.06203>

<https://arxiv.org/abs/2006.04757>

<https://arxiv.org/abs/2109.00110>

Current members

Bartosz Piotrowski, Jelle Piepenbrock, Josef Urban, Natasha Seelam, Debajyoti Datta, Leon Weber, Zsombor Koman, Pedro Ortiz Suárez, Zaid Alyafeai, Stella Biderman, Aashiq Muhamed, Rishabh Anand



Current chair(s)

Jelle Piepenbrock, Bartosz Piotrowski

Current status

In preliminary stages, planning and scoping.

14. Language models for historical texts

Scope

This working group focuses on the development and evaluation of language models for historical texts (in English and in other languages). Relevant topics include: selecting and preprocessing primary and secondary sources, tokenization and parsing, dealing with OCR and OCR errors, transfer learning between modern and historical language models, model training, and evaluation. Evaluation should also include tasks that are specific to historical datasets, such as semantic change, and tasks which are relevant to the GLAM (Galleries, Libraries, Archives and Museums) sector.

Resources - onboarding - documentation

<https://drive.google.com/drive/folders/1kNH8sXL2ZyhSdiHQvvgpU2JVDeaxUJjz>

Current members

Francesco De Toni, Javier de la Rosa, Daniel van Strien, Maraim Masoud, Abheesht Sharma, Zsombor Koman, Khalid Almubarak, Clémentine Fourier, Enrique Manjavacas, Isaac Rodríguez, Harrison Pim, Christopher Akiki, Jörg Frohberg, Lauren Klein, Huy Ngo, Ali Hürriyetoglu, Jordan Clive

Current chair(s)

Francesco De Toni, Daniel van Strien

Current status

In preliminary stages. First meeting to be organised soon to define scope, plan and expected output.



Proposal for New Working Groups

Here you can propose new Working Groups. Feel free to propose yourself as a Chair or just propose a Working Groups idea. You should fill a-minima the “Scope”.

Template for new Working Groups (feel free to copy it above and propose new WG)

Scope

Resources - onboarding - documentation

Current members

Current chair(s)

Current status



Code-mixed datasets and language models

Scope: Code-mixing is a frequent style of communication for multilingual speakers. The majority of the resources (e.g., multilingual language models) for multilingual research do not work effectively with the code-mixing tasks. In this WG, we plan to create the datasets and language model for code-mixed languages (e.g., Hindi-English, Spanish-English, Bengali-English, etc.). The currently available datasets for most of the code-mixing tasks such as NER, POS tagging, and machine translation are extremely small in size (<10,000 samples). We will explore ways to generate and curate datasets at a large scale.

Resources - onboarding - documentation

Current members

Current chair(s)

Current status