

Meeting of the LF AI & Data Technical Advisory Council (TAC)

April 7, 2022

 LF AI & DATA

Antitrust Policy

- › Linux Foundation meetings involve participation by industry competitors, and it is the intention of the Linux Foundation to conduct all of its activities in accordance with applicable antitrust and competition laws. It is therefore extremely important that attendees adhere to meeting agendas, and be aware of, and not participate in, any activities that are prohibited under applicable US state, federal or foreign antitrust and competition laws.
- › Examples of types of actions that are prohibited at Linux Foundation meetings and in connection with Linux Foundation activities are described in the Linux Foundation Antitrust Policy available at <http://www.linuxfoundation.org/antitrust-policy>. If you have questions about these matters, please contact your company counsel, or if you are a member of the Linux Foundation, feel free to contact Andrew Updegrove of the firm of Gesmer Undergone LLP, which provides legal counsel to the Linux Foundation.

Recording of Calls

Reminder:

TAC calls are recorded and available for viewing on the [TAC Wiki](#)

Reminder: LF AI & Data Useful Links

- › Web site: lfaidata.foundation
- › Wiki: wiki.lfaidata.foundation
- › GitHub: github.com/lfaidata
- › Landscape: <https://landscape.lfaidata.foundation> or <https://l.lfaidata.foundation>
- › Mail Lists: <https://lists.lfaidata.foundation>
- › Slack: <https://slack.lfaidata.foundation>
- › Youtube: <https://www.youtube.com/channel/UCfasaeqXJBCAJMNO9HcHfbA>
- › LF AI Logos: <https://github.com/lfaidata/artwork/tree/master/lfaidata>
- › LF AI Presentation Template: https://drive.google.com/file/d/1eiDNJvXCqSZHT4Zk_-czASlz2GTBRZk2/view?usp=sharing

- › Events Page on LF AI Website: <https://lfaidata.foundation/events/>
- › Events Calendar on LF AI Wiki (subscribe available): <https://wiki.lfaidata.foundation/pages/viewpage.action?pageId=12091544>
- › Event Wiki Pages: <https://wiki.lfaidata.foundation/display/DL/LF+AI+Data+Foundation+Events>

Agenda

- › Roll Call (2 mins)
- › Approval of Minutes from previous meeting (2 mins)
- › RosaeNLG annual review (15 minutes)
- › 1chipML - Free/Open library for Math and ML on microcontrollers (40 minutes)
- › LF AI General Updates (2 min)
- › Open Discussion (2 min)

TAC Voting Members - Please note

Please ensure that you do the following to facilitate smooth procedural quorum and voting processes:

- Change your Zoom display name to include your First/Last Name, Company/Project Represented
 - example: Nancy Rausch, SAS
- State your First/Last Name and Company/Project when submitting a motion
 - example: First motion, Nancy Rausch/SAS

TAC Voting Members

Note: we still need a few designated backups specified on [wiki](#)

Member Company or Graduated Project	Membership Level or Project Level	Voting Eligibility	Country	TAC Representative	Designated TAC Representative Alternates
4paradigm	Premier	Voting Member	China	Zhongyi Tan	
Baidu	Premier	Voting Member	China	Ti Zhou	Daxiang Dong, Yanjun Ma
Ericsson	Premier	Voting Member	Sweden	Rani Yadav-Ranjan	
Huawei	Premier	Voting Member	China	Howard (Huang Zhipeng)	Charlotte (Xiaoman Hu) , Leon (Hui Wang)
IBM	Premier	Voting Member	USA	Susan Malaika	Saishruthi Swaminathan
Nokia	Premier	Voting Member	Finland	@ Michael Rooke	@ Jonne Soininen
OPPO	Premier	Voting Member	China	Jimin Jia	
SAS	Premier	Voting Member	USA	*Nancy Rausch	JP Trawinski
Tencent	Premier	Voting Member	China	Bruce Tao	Huaming Rao
ZTE	Premier	Voting Member	China	Wei Meng	Liya Yuan
Acumos Project	Graduated Technical Project	Voting Member	USA	Amit Kumar	Prasanna Kulkarni
Adversarial Robustness Toolbox Project	Graduated Technical Project	Voting Member	USA	Beat Buesser	
Angel Project	Graduated Technical Project	Voting Member	China	Bruce Tao	Huaming Rao
Egeria Project	Graduated Technical Project	Voting Member	UK	Mandy Chessell	Nigel Jones, David Radley, Maryna Strelchuk, Ljupcho Palashevski, Chris Grote
Flyte Project	Graduated Technical Project	Voting Member	USA	Ketan Umare	
Horovod Project	Graduated Technical Project	Voting Member	USA	Travis Addair	
Milvus Project	Graduated Technical Project	Voting Member	China	Xiaofan Luan	Jun Gu
ONNX Project	Graduated Technical Project	Voting Member	USA	Alexandre Eichenberger	Prasanth Pulavarthi, Jim Spohrer
Pyro Project	Graduated Technical Project	Voting Member	USA	Fritz Obermeyer	

Minutes approval


Approval of March 24, 2022 Minutes

Draft minutes from the March 24th TAC call were previously distributed to the TAC members via the mailing list

Proposed Resolution:

- › That the minutes of the March 24 meeting of the Technical Advisory Council of the LF AI & Data Foundation are hereby approved.

Annual Review for RosaeNLG




04/07/2022

DLF AI & DATA

Sandbox Project review criteria

To be accepted into the Incubation stage, a project must meet all the requirements of the Sandbox stage plus:

<ul style="list-style-type: none"> ● Have at least two organizations actively contributing to the project. 	<p>No. Addventa is starting to contribute, along with the original author.</p>
<ul style="list-style-type: none"> ● Have a defined Technical Steering Committee (TSC) with a chairperson identified, with open and transparent communication. 	<p>Not in place</p>
<ul style="list-style-type: none"> ● Have a sponsor who is an existing LF AI & Data member. Alternatively, a new organization would join LF AI & Data and sponsor the project's incubation application. 	<p>Orange is the original sponsor</p>
<ul style="list-style-type: none"> ● Have at least 300 stars on GitHub; this is an existing requirement for a project to be listed on the LF AI & Data landscape. 	<p>No. 52 stars in March 2022.</p>
<ul style="list-style-type: none"> ● Have achieved and maintained a Core Infrastructure Initiative Best Practices Silver Badge. 	<p>Yes. Effort can be improved on continuity of maintenance and on security.</p>
<ul style="list-style-type: none"> ●  In addition to the affirmative vote of the TAC, incubation stage projects also require the affirmative vote of the Governing Board. 	<p>NA - project did not request to upgrade to Incubate</p>

RosaeNLG



Brief Description:

RosaeNLG is a template-based Natural Language Generation (NLG) automates the production of relatively repetitive texts based on structured input data and textual templates, run by a NLG engine. Production usage is widespread in large corporations, especially in the financial industry.

Typical use cases are:

- describing a product based on its features for SEO purposes
- produce structured reports like risk reports or fund performance in the financial industry
- generate well formed chatbot answers.

RosaeNLG is an open source NLG engine. It aims:

- to offer the same NLG features as product NLG solutions
- to be developer and IT friendly for template configuration and integration
- to provide NLG on both server-side and browser-side

Contributed by:

Ludan Stoecklé in March 2021 as a Sandbox Project

Key Links:

Github: <https://github.com/RosaeNLG/rosaenlg>



Website: <https://rosaenlg.org>

Artwork:

<https://github.com/lfai/artwork/tree/master/projects/rosaenlg>

Mailing lists:

- › [rosaenlg-announce](#)
- › [rosaenlg-technical-discuss](#)
- › [rosaenlg-tsc](#)



Organizations contributing



Ludan Stoecklé, the original author, remains the main contributor.




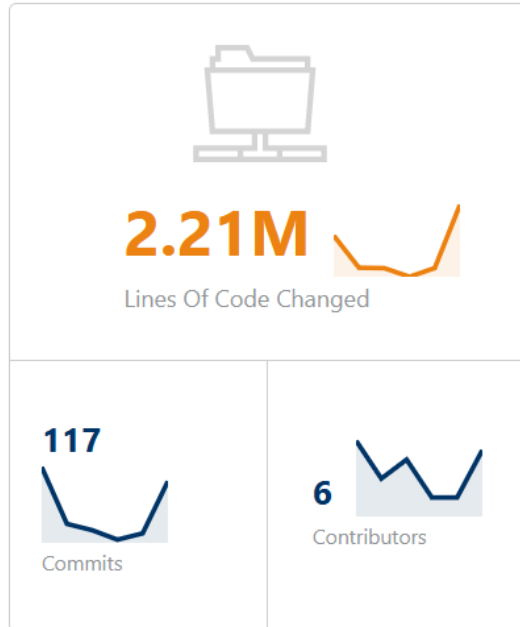
Addventa (AI company based in Paris, France) is reporting bugs regularly, suggesting features regularly, and contributing to code (e.g. #97, #116).

Contributions

SOURCE CONTROL

Commits

[Go To Overview](#)



Top 10 Contributors By Commits

[View All](#)

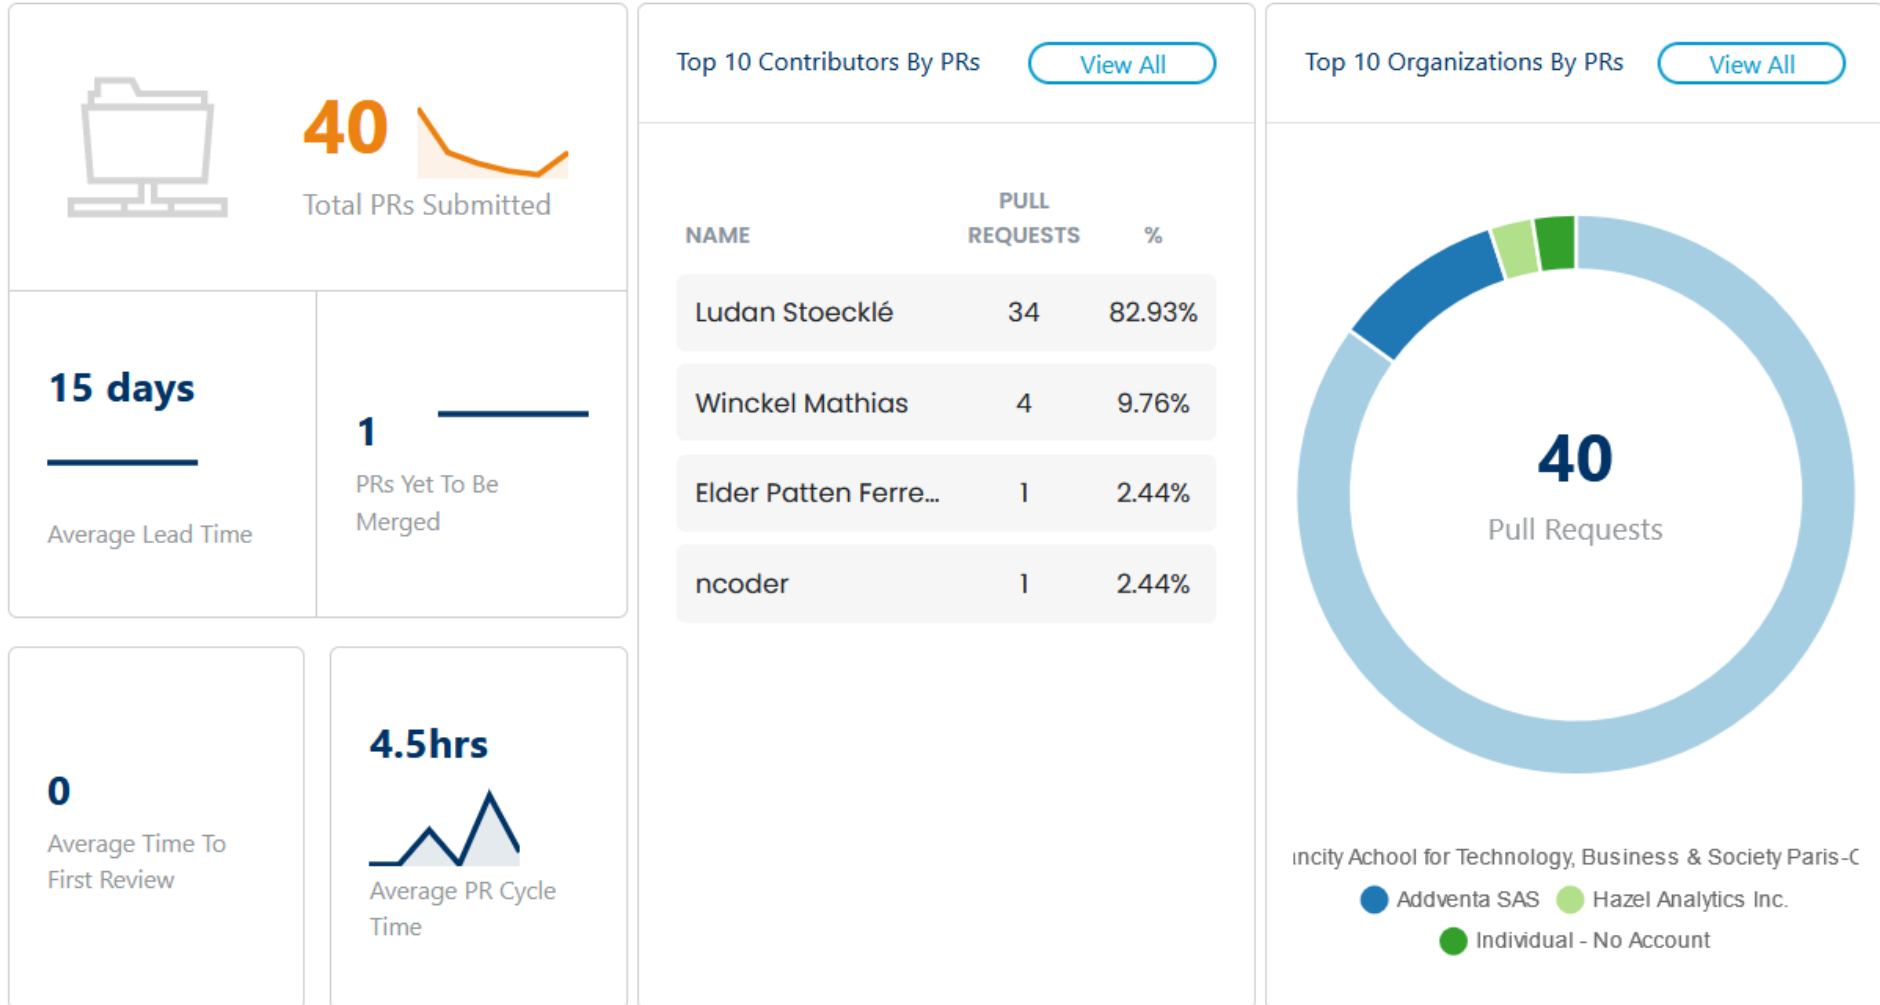
NAME	LINES OF CODE	COMMITTS
Ludan Stoecklé	2.21M	106
Ibrahim Haddad	31	5
Winckel Mathias	92	3
Elder Patten Ferre...	2	1
ncoder	5	1
pY6xJKN8pN	13	1

Top 10 Organizations By Commits

[View All](#)



Contributions



Key Achievements in the past year

- › “routine maintenance” is done: bug fixes, libraries security upgrades (e.g. Log4Shell), maintenance of documentation, with 16 versions released in 2021
- › the users are reporting bugs (#47, #55, #112), do feature requests (#72, #84, #109), and are starting to code (#97, #116)
- › the users are showing a high expertise in their bug reports and feature requests
- › 3 new features released, all suggested by end users:
 - › usage without templates (low level API)
 - › 2 levels when assembling sentences dynamically
 - › “universal tenses”, which work for multiple languages

Areas the project could use help on


- › The project needs more ideas to feed the roadmap and more real contributors.
- › Find a way to incentive Addventa to contribute more intensively.

Feedback on working with LF AI & Data

- › LF AI & Data has brought long term **trust** and **visibility** for corporate usage.
- › The enrolling process (sandbox criteria) in itself raises the quality bar of the code and its documentation.



Free/Open library for Math and ML on microcontrollers



Jean Michel Sellier,
GAIA Ericsson,
Montreal

OLFAI & DATA



Initial note

- › The material presented in these slides is partially extracted from our formal proposal.
- › This proposal was submitted via GitHub and can be found at the following URL:


<https://github.com/lfai/proposing-projects/blob/master/proposals/1chipML.md>





Topics

- › *Introduction to various aspects of the library*



- › *1chipML and LF AI & Data*



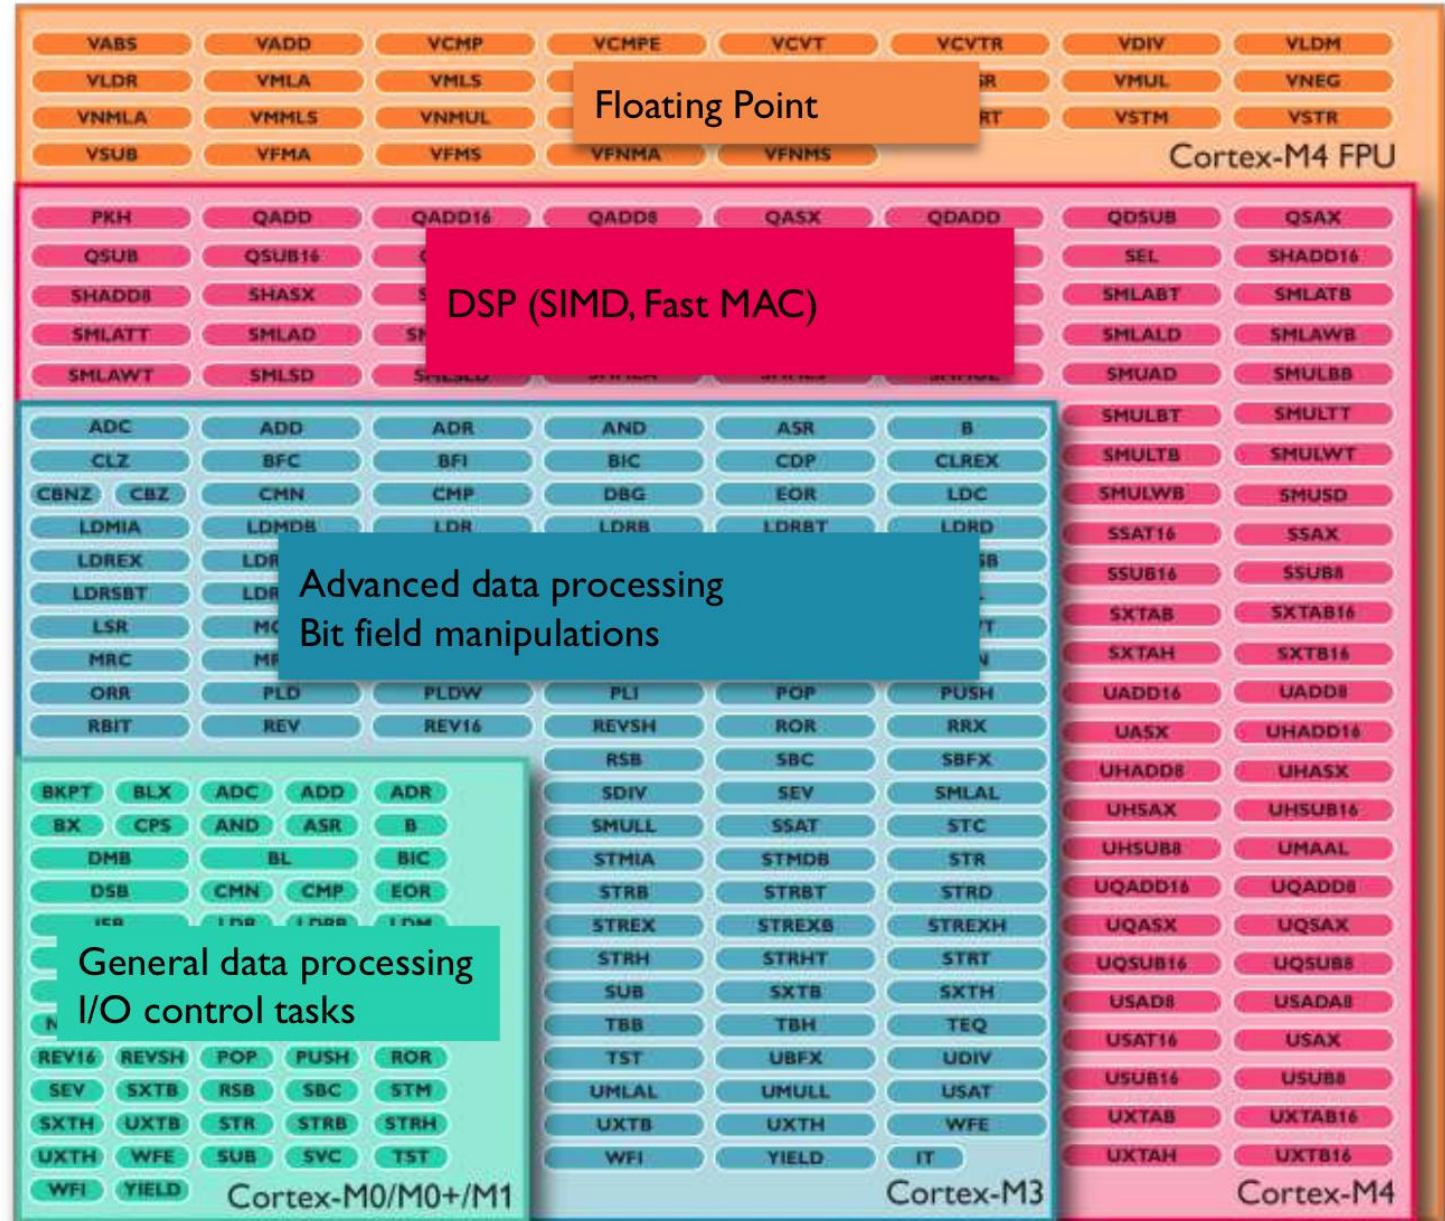
ARM Cortex-M Instruction Set Architecture



Cortex-M4

Cortex-M3

Cortex-M0/M0+





Various ways to develop code for MCUs

From Arduino IDE



```
sketch_dec07a | Arduino 1.8.3  
File Edit Sketch Tools Help  
sketch_dec07a  
void setup() {  
  // put your setup code here, to run once:  
  
}  
  
void loop() {  
  // put your main code here, to run repeatedly:  
  
}
```

2 Arduino/Genuino Uno on COM3

From the shell



```
Terminal - sellier@sellier-VirtualBox: ~  
File Edit View Terminal Tabs Help  
(base) sellier@sellier-VirtualBox:~$
```



How to embed 1chipML?

- › The whole library can be included easily in any program and/or environment by simply adding:

```
#include<1chipml.h>
```

- › An alternative option is to include only the methods that are needed:

```
#include<1chipml/gauss_elimination.h>
```

In this case, the library is structured in such a way that only the necessary routines will be embedded into the user code.



A concrete example

```
int main(void){
    // ...
    gauss_real A[2][2];
    gauss_real B[2];
    gauss_real *sol; /* pointer towards the solution */

    /* system of 2 linear equations A*X=B */
    N=2;
    A[0][0]=+1.; A[0][1]=+1.;
    A[1][0]=+1.; A[1][1]=-1.;
    B[0]=0.;
    B[1]=1.;

    /* Apply gauss elimination method to solve the system */
    sol=gauss_elimination(N,A,B);
    for(i=0;i<N;i++) printf("sol[%d] = %0.3f\n",i,sol[i]);
    return(0);
}
```



Basic number crunching

Source: “*Numerical Recipes*”, S. Teukolsky, W.H. Press, W.T. Vetterling, Cambridge University Press, (2007).

- › Linear systems:
Gaussian elimination (a.k.a. row reduction), LU decomposition.
- › Random number generation:
Linear congruential generator, Mersenne twister.
- › Interpolation/Extrapolation:
Polynomial and spline interpolation/extrapolation.
- › Optimization problems:
Gradient descent method, genetic algorithm.
- › Function numerical derivation and integration:
First and second-order finite differences approaches, Monte Carlo approach.
- › Eigenproblems:
Jacobi method, Lanczos method.



Statistical methods

Source: Wikipedia pages and “Numerical Recipes” (see previous slide).

- › Analysis of variance
Naïve approach, shifted data variance.
- › Cluster analysis
k-means clustering, distribution-based clustering, grid-based clustering.
- › Correlation
Naïve application of the definition.
- › Curve fitting
Same as in previous slide.
- › Regression
Linear and multi-linear regression methods.



Neural networks in Tensorflow, or PyTorch

Source: https://www.tensorflow.org/tutorials/keras/save_and_load

- › Both Tensorflow and Pytorch offer the possibility of saving an ANNs after the training phase. The ANNs can, consequently, be loaded by other

```
# Save the weights
model.save_weights('./checkpoints/my_checkpoint')

# Create a new model instance
model = create_model()

# Restore the weights
model.load_weights('./checkpoints/my_checkpoint')

# Evaluate the model
loss, acc = model.evaluate(test_images, test_labels, verbose=2)
print("Restored model, accuracy: {:.2f}%".format(100 * acc))
```



Fast Fourier Transform (FFT)

Source: Wikipedia

› Cooley-Tukey algorithm

```
 $X_{0,\dots,N-1} \leftarrow \text{ditfft2}(x, N, s):$   
  if  $N = 1$  then  
     $X_0 \leftarrow x_0$   
  else  
     $X_{0,\dots,N/2-1} \leftarrow \text{ditfft2}(x, N/2, 2s)$   
     $X_{N/2,\dots,N-1} \leftarrow \text{ditfft2}(x+s, N/2, 2s)$   
    for  $k = 0$  to  $N/2-1$  do  
       $p \leftarrow X_k$   
       $q \leftarrow \exp(-2\pi i/N k) X_{k+N/2}$   
       $X_k \leftarrow p + q$   
       $X_{k+N/2} \leftarrow p - q$   
    end for  
  end if
```

DFT of $(x_0, x_s, x_{2s}, \dots, x_{(N-1)s})$:

trivial size-1 DFT base case

DFT of $(x_0, x_{2s}, x_{4s}, \dots)$

DFT of $(x_s, x_{s+2s}, x_{s+4s}, \dots)$

combine DFTs of two halves into full DFT:



Reinforcement learning

Source: “*Reinforcement Learning: An Introduction*”, R.S. Sutton, A.G. Barto, MIT Press, (2015).

- › Multi-arm bandits problem:
 ϵ -greedy method,
UCB (upper confidence bound) method,
gradient bandits method.
- › Monte Carlo methods:
On-policy methods,
Off-policy methods.




Practical prototypes

- › Once we release a major version of the library, we plan to regularly create a few prototypes to show how the library can be used in practical cases.

This should be twofold:

It provides further documented examples which can be utilized by the users.


It helps increasing the community around the library because practical examples are a way to clearly show the capabilities of the library.




















The GitHub repository

<https://github.com/1chipML/1chipML>



 jmsellier Update GOVERNANCE.md	f58fc5c 19 hours ago	 20 commits
 doc	Delete empty.txt	13 days ago
 src	removed dummy file	20 days ago
 tests	remove dummy file	20 days ago
 CODE_OF_CONDUCT.md	Update CODE_OF_CONDUCT.md	19 hours ago
 CONTRIBUTING.md	Update CONTRIBUTING.md	19 hours ago
 GOVERNANCE.md	Update GOVERNANCE.md	19 hours ago
 LICENSE	first files added	20 days ago
 MAINTAINERS.md	first files added	20 days ago
 Makefile	first files added	20 days ago
 README.md	Update README.md	19 hours ago
 RELEASES.md	first files added	20 days ago
 SECURITY.md	Add files via upload	13 days ago
 SUPPORT.md	Add files via upload	13 days ago



The team




- › The current 1chipML team has long and thorough experiences with numerical crunching on different platforms ranging from simple microcontrollers to huge supercomputers.
- › The team is currently made of data scientist at GAIA, Ericsson (Global AI Accelerator) with relevant skills in machine learning, C programming, parallel algorithms, edge computing, as well as experience with embedded systems, open-source codes and.. a lot of enthusiasm!



JM Sellier




Q. Qi



S. Memon



H. Dalal



Y. Nie



S. Bashbaghi



P. Karthikeyan



G. Kunz



Why LF AI & Data?

- › There are plenty of free/open communities, but LF AI & Data is practically the only authoritative community which gathers around the concepts of artificial intelligence and data science.
- › Some of us has past experience in developing, spreading and maintaining GNU packages. While this experience is extremely constructive and valuable, GNU does not provide any particular and concrete way to support the development of a community.. and this is exactly one of the main goals we have with 1chipML!




How does 1chipML fit into LF AI & Data mission?

- › **From the charter:**

...The primary mission of the LF AI & Data Foundation is to raise, budget and spend funds in support of various artificial intelligence, machine learning, and data-related open source projects, including infrastructure and support initiatives related thereto...

- › **Our mission with 1chipML:**


Create a community around our Free/Open source library for the continuous development and improvement of a library for basic math and machine learning for microcontrollers.





What unique value this project provides to LF AI, its community of companies and developers?

- › *The targeted hardware in 1chipML is AVR/ARM microcontrollers. The idea is to enable concrete and practical Internet of Things and Edge Computing on such low power, limited resources computing devices.*
- › *From an inspection of the page <https://lfaidata.foundation/projects/>, it seems that there is no project similar to 1chipML in the LF AI community, as no other project specifically seems to target microcontrollers.*





What future impact this project could have for the industry and AI community through LF AI?

› **Internet of Things (IoT)**


The IoT describes the network of physical objects that are embedded with sensors, software, and other technologies that is used for the purpose of connecting and exchanging data with other devices and systems over the Internet.

...A challenge for producers of IoT applications is to clean, process and interpret the vast amount of data which is gathered by the sensors...



› **Edge computing**


The objectives of Distributed Artificial Intelligence are to solve the reasoning, planning, learning and perception problems of artificial intelligence, especially if they require large data, by distributing the problem to autonomous processing nodes.





License, tools, dependencies, etc.

- › License: Apache 2.0
- › Tools: gcc, AVR and ARM versions of gcc specific for their microcontrollers (free/open)
- › Issue tracker: GitHub tracker
- › Collaboration tools: GitHub requests, mailing lists, wiki pages, Slack channel
- › External dependencies: none






Contributors, committers, maintainers, etc.



- › **Roles of contributors, committers, maintainers, etc. defined in:**
<https://github.com/1chipML/1chipML/blob/main/MAINTAINERS.md>
[1chipML/CONTRIBUTING.md at main · 1chipML/1chipML \(github.com\)](https://github.com/1chipML/1chipML/blob/main/CONTRIBUTING.md)
- › **Contributors** are community members who contribute in concrete ways to the project. Anyone can become a contributor, and contributions can take many forms... For example:
 - supporting new users,
 - commenting on issues,
 - code reviewing,
 - writing, editing, translating or reviewing the documentation,
 - organizing events or evangelizing the project, ...and much more!
- › **Contributing:** As a contributor, you'll have to follow the rules: When contributing to 1chipML, please first talk to us, then make or share your suggested change by pulling a request on the GitHub 1chipML page. Once your proposition has been accepted, you will be able to start your contribution. Once the contribution is completed, you can pull a new request so that the maintainers can validate and share your work with the next library release.



Release methodology



- › **Release methodology defined in:**
<https://github.com/1chipML/1chipML/blob/main/RELEASES.md>

- › **Methodology:** *The main guiding principle in releasing a new version is always: the quality of the master branch should never be compromised.*

When a coder feels ready to share a contribution with the community, a pull request is sent to the maintainers.

We plan to review every pull request directly from the GitHub website. This allows public commentaries and provides transparency for all users. While sharing feedback always be reasonable, civil, courteous, and kind. Disagreement is obviously acceptable, but the conversation has to be polite and respectful.

- › **Questions:** *Do this changes make sense? Does this change have positive impact on the library? Is this change really a new feature?*

- › **Merge approval:** The project maintainers include a GFM (Good For Me) in the comments on the pull request to indicate acceptance prior to merging. A change requires "GFMs" from two project maintainers. In the particular case that the code is written by a maintainer, the change only requires one additional "GFM".



Code of conduct



- › **Code of conduct defined:**
https://github.com/1chipML/1chipML/blob/main/CODE_OF_CONDUCT.md
- › **Verbatim:** *“The 1chipML community is dedicated to providing a harassment-free experience for everyone, regardless of gender, gender identity and expression, sexual orientation, disability, physical appearance, body size, age, race, or religion. We do not tolerate harassment of participants in any form. This code of conduct applies to all 1chipML spaces. Anyone who violates this code of conduct may be sanctioned or expelled from these spaces at the discretion of the 1chipML maintainers team. All participants are responsible for knowing and abiding by these rules.”*
- › **Reporting code of conduct issues:** internally, when possible, or conduct@lfai.foundation.



Project governance



- › **Project governance defined:**
<https://github.com/1chipML/1chipML/blob/main/GOVERNANCE.md>
- › **Verbatim:** *“The 1chipML project aims for open and transparent governance and decision-making... 1chipML is a meritocratic, consensus based project. Anyone interested in joining this community is welcome and encouraged to contribute...”*
- › Also, refer to “contributing”, “code_of_conduct” and “maintainers” MD files.



Possible collaborations with other LF AI projects

- › ONNX; this project offers an open standard to represent machine learning models and provides ways to maximize performances across hardware.
- › Horovod; this project provides a distributed deep learning training framework for known ML libraries/frameworks such Tensorflow, etc.
- › Angel ML; this is a high-performance distributed machine learning platform.





Thanks!

Upcoming TAC Meetings

 **DLF** AI & DATA

Upcoming TAC Meetings

- › April 21, 2022 – Devops Committee update, Outreach Committee update
- › May 5, 2022 – Open

Please note we are requesting special topics for future meetings.

If you have a topic idea or agenda item, please send agenda topic requests to tac-general@lists.lfaidata.foundation

Open Discussion

TAC Meeting Details

- › To subscribe to the TAC Group Calendar, visit the wiki:
<https://wiki.lfaidata.foundation/x/cQB2> _____
- › Join from PC, Mac, Linux, iOS or Android: <https://zoom.us/j/430697670>
- › Or iPhone one-tap:
 - › US: +16465588656,,430697670# or +16699006833,,430697670#
- › Or Telephone:
 - › Dial(for higher quality, dial a number based on your current location):
 - › US: +1 646 558 8656 or +1 669 900 6833 or +1 855 880 1246 (Toll Free) or +1 877 369 0926 (Toll Free)
- › Meeting ID: 430 697 670
- › International numbers available: <https://zoom.us/u/achYtcw7uN>

Legal Notice

- › The Linux Foundation, The Linux Foundation logos, and other marks that may be used herein are owned by The Linux Foundation or its affiliated entities, and are subject to The Linux Foundation's Trademark Usage Policy at <https://www.linuxfoundation.org/trademark-usage>, as may be modified from time to time.
- › Linux is a registered trademark of Linus Torvalds. Please see the Linux Mark Institute's trademark usage page at <https://lmi.linuxfoundation.org> for details regarding use of this trademark.
- › Some marks that may be used herein are owned by projects operating as separately incorporated entities managed by The Linux Foundation, and have their own trademarks, policies and usage guidelines.
- › TWITTER, TWEET, RETWEET and the Twitter logo are trademarks of Twitter, Inc. or its affiliates.
- › Facebook and the "f" logo are trademarks of Facebook or its affiliates.
- › LinkedIn, the LinkedIn logo, the IN logo and InMail are registered trademarks or trademarks of LinkedIn Corporation and its affiliates in the United States and/or other countries.
- › YouTube and the YouTube icon are trademarks of YouTube or its affiliates.
- › All other trademarks are the property of their respective owners. Use of such marks herein does not represent affiliation with or authorization, sponsorship or approval by such owners unless otherwise expressly specified.
- › The Linux Foundation is subject to other policies, including without limitation its Privacy Policy at <https://www.linuxfoundation.org/privacy> and its Antitrust Policy at <https://www.linuxfoundation.org/antitrust-policy>. each as may be modified from time to time. More information about The Linux Foundation's policies is available at <https://www.linuxfoundation.org>.
- › Please email legal@linuxfoundation.org with any questions about The Linux Foundation's policies or the notices set forth on this slide.