

VOLUME 41
NUMBER 1

I/O VIVAT

From Idea to Market

Reward-based Crowdfunding for High-tech Gadgets

Computer Choreography

Dutch Education System

Why the Netherlands is so unique

Separation is Hard

Using Linux Namespaces

And more...

Committees & Columns

Guest writers



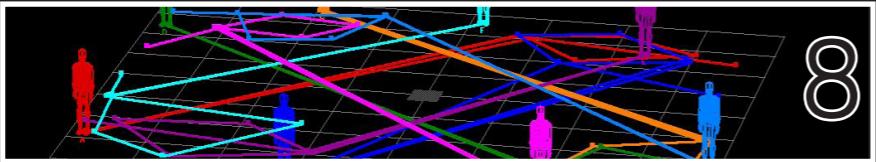
Inter-Actief



Reward-based Crowdfunding for High-tech Gadgets



GenAI Masterclass



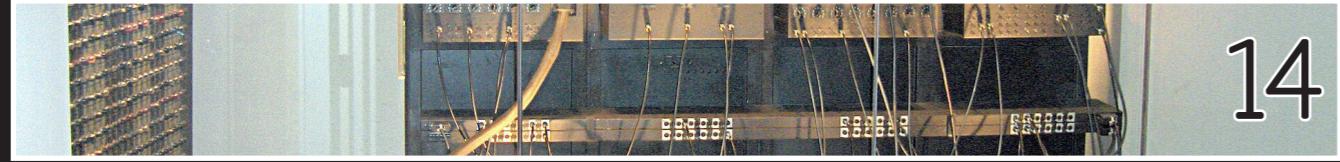
Computer Choreography



From the Chairwoman



Linux Namespaces



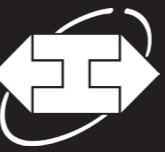
From the ENIAC Board



Dutch Education System



Hyper-Actief



Volume 41, number 1
November 2025
ISSN: 1389-0468

I/O Vivat is the popular scientific magazine of I.C.T.S.V. Inter-Actief, the study association for Technical Computer Science, Business Information Technology and the corresponding Master's programmes at the University of Twente.

// Chief Editors

Anamaria Ceban
Erjan Steenbergen

// Editors

Filip Karakalasev, Jeske van Hijum,
Ruben Groot Roessink

// Guest Writers

Claire van Teunenbroek,
Pauline Weritz, Johannes Dahlke,
Koen Braham, Kevin Alberts,
Tom Grooters, Josje van 't Padje

For questions,
comments or suggestions,
I/O Vivat can be reached via
e-mail at iovivat@inter-actief.net,
by phone via 053-489 3756 or
by mail:
Study association Inter-Actief
PO box 217, 7500AE Enschede

© 2025 I.C.T.S.V. Inter-Actief



I/O VIVAT

//Editorial

Dear reader,

We are back for the 41st volume of the I/O Vivat! This academic year, we start off with another digital version of our I/O Vivat. Due to budgetary limits and absence of company partners, the physical version this year is postponed to a slightly later date while we work on finding more partners to keep sustaining the ability to physically print the I/O Vivat. In the meantime, you can enjoy the new articles that have been written in the summer break by our lovely editors and guest writers! Also, do not miss our blogposts on our digital platform (<https://ioviv.at>). Do not hesitate to scan the QR code down below to sign up for our website, so that you will not miss any articles in the future!

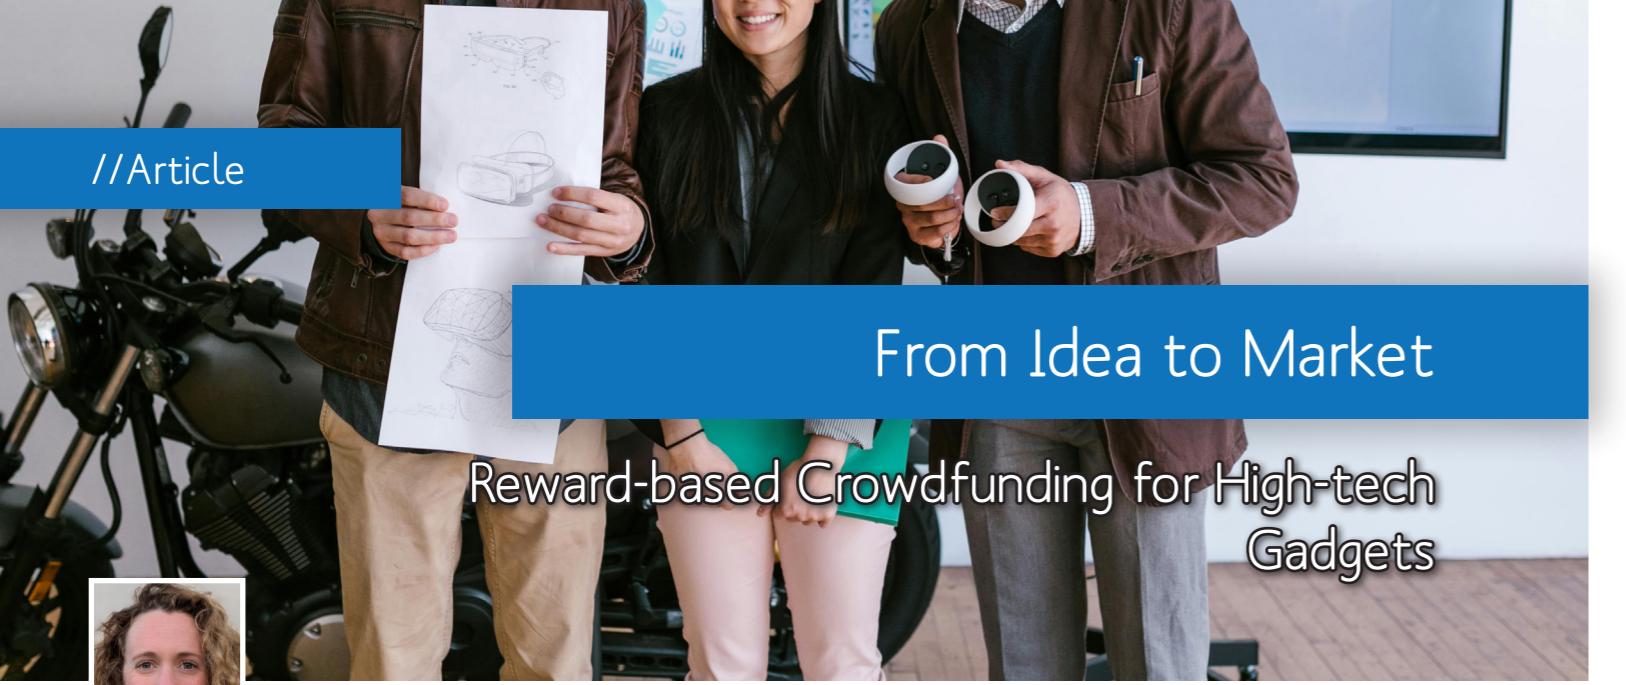


In this bite-sized edition, we explore the ins and outs of reward-based crowdfunding, how AI extends further from solely computer science, take a closer look on the intricacies of the Dutch education system, receive a masterclass on solving CTF challenges using Linux namespaces, and find the perfect balance between computer and dance. Furthermore, we have the honour to once again give you columns from both the board of Inter-Actief and ENIAC, as well as the Hyper-Actief committee!

If you are reading this and think to yourself: "I wish I could write an article in the I/O Vivat," do not hesitate to reach us via e-mail (iovivat@inter-actief.net). We are happy to invite anyone who is eager or interested to write! Perhaps you will be able to find yourself on the next edition!

We thank you for reading the I/O Vivat!

Anamaria Ceban & Erjan Steenbergen
Chief Editors I/O Vivat



By: Claire van Teunenbroek
Assistant Professor of Entrepreneurship & Marketing UT

Tech projects within tech start-ups are at the centre of science-driven and technology-based innovation, aiming to address complex societal problems by developing cutting-edge new technologies. However, the factors that drive these tech start-up projects—multidisciplinary research, technical diversity, complexity, and prolonged development timelines—make them inherently risky and challenging to finance. Innovative financial solutions like reward-based crowdfunding have emerged as complements to government financing, institutional and private investors' funding [1].

Reward-based crowdfunding, henceforth referred to as 'crowdfunding', has been around for a quarter of a century. It's an online funding mechanism used by individuals to promote a product still in development. As such, it's an in-

teresting option for projects nearing the product development stage, since it offers a viable alternative to conventional bank loans and equity capital [2]. What exactly is crowdfunding, and what can it do in terms of funding product development? How many products reach the market, and what can explain successful market entry?

We discuss crowdfunding characteristics, its place in the funding trajectory and the three crowdfunding stages. Moreover, we outline failure rates and how to navigate the crowdfunding process.

Initiators, backers, and platforms

Via crowdfunding, entrepreneurs (i.e. initiators) appeal to a group of individuals (i.e. backers) by presenting their concepts on a digital platform (i.e. a crowdfunding platform). Platforms like *Kickstarter* and *Indiegogo* are popular

examples of such platforms. Start-ups and individuals commonly use crowdfunding to validate ideas, build a community, and raise funds without giving up equity. As one entrepreneur put it [3]:

"We were only in our mid/late twenties and we didn't want to immerse ourselves in debt or anything like that. So crowdfunding felt more noncommittal."

In return for a financial contribution, backers can receive rewards, ranging from a simple thank you to pre-ordering the product.

In 2024, the global reward-based crowdfunding market was valued at approximately \$7.45 billion, representing 50% of the total crowdfunding market, which stood at \$14.89 billion [5]. *Kickstarter* is the leading platform for this.

Stage I: Pre-campaign stage

Ideation to campaign launch

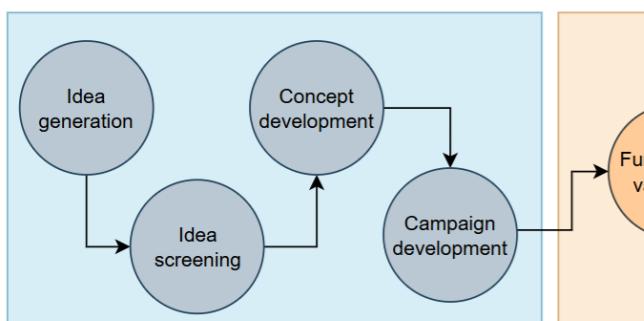
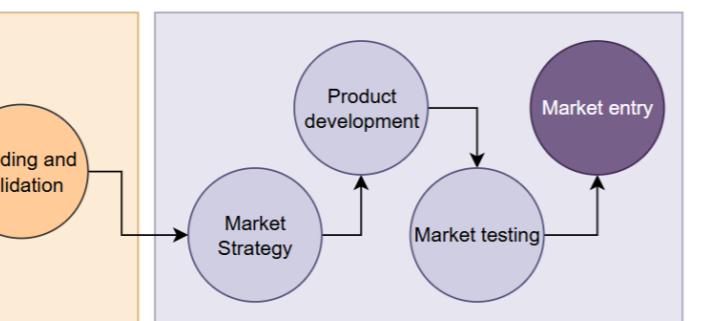


Figure 1: A simplified visualisation of the crowdfunding process

Stage II: Campaign stage

Live campaign stage



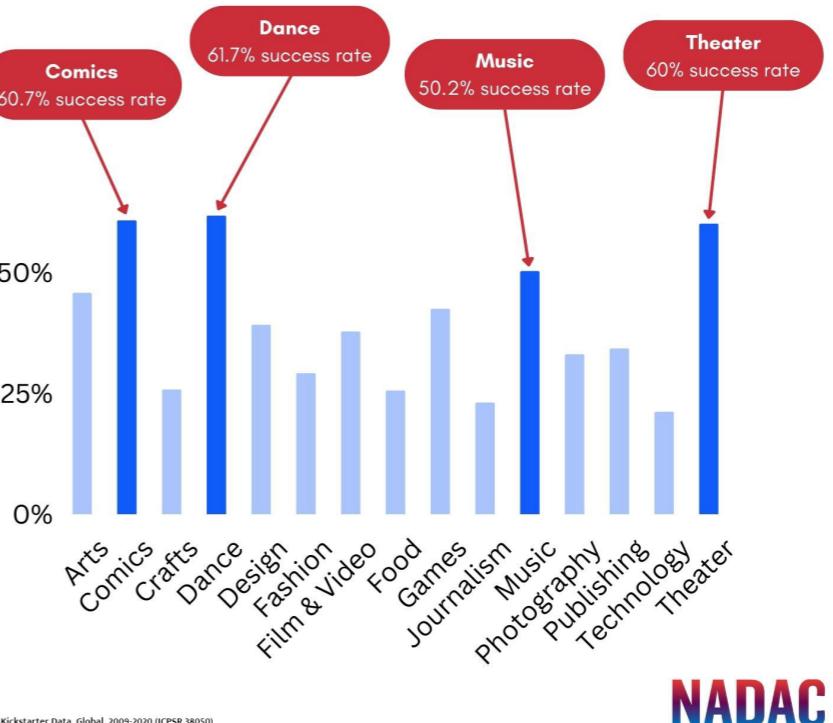
Stage III: Post-campaign stage

Post-campaign product delivery



Figure 2: Kickstarter success rates from various subcategories [6]

Arts-related Kickstarter Campaign Success Rates, 2009-2020



Source: Kickstarter Data, Global, 2009-2020 (ICPSR 38050)

NADAC

remain underexplored.

The Technology category is one of *Kickstarter*'s largest, with over \$1.46 billion pledged to more than 56,000 projects as of early 2025 [5]. The Technology category is third in total funding, behind Games and Design. In the Technology category, 'gadgets' are popular, like the project *See.Sense ICON*, which aims to develop a bike light that also functions as a warning system: it flashes next to emitting a constant light when a car approaches the bike. Another example is *Cyndr Sentry*, a motion detection device that alerts gamers wearing noise-cancelling headphones or VR gear when someone enters their space, using real-time visual or in-game notifications.

From beginning to end?

While crowdfunding is often framed as exclusively a funding strategy, it is important to approach it as a multi-stage process. The crowdfunding process includes three stages: (1) ideation to campaign launch, (2) the crowdfunding campaign, and (3) post-campaign product delivery (see Figure 1). Thus, collecting funds – Stage 2 – is only one component of the broader crowdfunding process. Interestingly, most academic attention focuses on Stage 2, while the beginning and end of the process

"Crowdfunding is a smart tool in the way that you can test your concept before ordering 10,000 pieces, you know."

The campaign stage is considered successful if the target amount is reached—

only then does the process continue into the post-campaign stage. While we know a good amount about factors impacting the success of collecting funds via crowdfunding, little is known about what happens once the funding has been collected. This is problematic since steps like market strategy, product development, and market testing are complex tasks that must be completed before a product can enter the market (see the post-campaign stage in Figure 1).

A high 'failure' rate

The funding stage often fails which means that a campaign did not collect the needed amount. *Kickstarter*, the world's leading reward-based crowdfunding platform, reports a failure rate of 58% [4], but this differs per category. For instance, the failure rate is higher in the Technology category, namely 76%, while it is 39% in the Dance category [4]. In the subcategory of High-Tech Gadgets, a failure rate of 67% was reported among projects run between 2012 and 2023 [6].

Recent findings focused on high-tech gadgets suggest that market entry is not guaranteed even after a successful crowdfunding campaign: only 43% of the projects that completed the campaign stage ultimately resulted in products available for sale online [7]. Given that only 33% of high-tech gadget projects on Kickstarter succeed in the campaign stage [6], the overall market entry rate of the High-Tech Gadgets category across both stages is approximately 14%. Notably, this does not represent the success rate of the entire crowdfunding process, as it is a maximum estimate. The actual rate is likely lower when accounting for the pre-campaign stage (Stage I), for which the success rate remains unknown.

We do not view failure to enter the market as a reason to diminish the value of crowdfunding, which serves purposes beyond financial acquisition [2]. The idea of 'success' in the context of crowdfunding is more complex than most academic studies suggest. The prevailing emphasis on reaching the funding target during the campaign stage equates crowdfunding success with capital raised. However, crowdfunding is also valuable in generating validation and

public awareness [2]. As one entrepreneur mentioned [3]:

"So the main idea behind it [crowdfunding] was of course to raise some money, but just as much to use it as a marketing platform, so to speak. So I could reach people in a different way than I would have otherwise."

Reach a large group of backers

What increases the chance of running a successful crowdfunding campaign? The answer differs per crowdfunding stage, but overall, the consensus is that it's important to reach a large group of backers [7,8]. Success in the post-campaign stage, specifically market entry, depends more on overall support – measured by total funds raised and the number of backers – than on the average contribution per backer. As such, we recommend that initiators extend their outreach beyond their immediate network to engage as many potential backers as possible. Although one might expect this to be particularly relevant for projects with higher target amounts, findings contradict this assumption: small and large projects benefit from attracting a large group of backers [7].

To attract a large group of backers, research suggests focusing on three key elements: social, quality, and rewards [8]. Social aspects refer to the initiator's communication skills and personal network. Early backers are often friends, family, or acquaintances, so having a large and active network is crucial. In the campaign's initial phase,



strong communication skills help convince individuals to support the project. Quality signals become more important once the initiator's network is tapped out. These signals can take many forms, such as the clarity and tone of the project description. An overly technical text may alienate potential backers, while a well-written description in layman's terms—balanced with enough detail to convey the product's value—can be far more persuasive.

Finally, social proof can signal quality. The number of backers in the early days can influence others to join. The behaviour of others often affects people, so campaigns that gain momentum quickly are more likely to continue attracting support.

In reward-based crowdfunding, the rewards are central—they're often the primary reason backers choose to contribute. While tangible rewards like early access to a product are important, successful campaigns also tap into immaterial value: the sense of exclusivity, involvement, and experience. For example, offering decision power—such as letting backers vote on design features or choose between colour options—can foster a deeper connection between the backer and the project. These kinds of rewards don't just deliver a product; they deliver a personalised experience that makes backers feel part of the journey.

In short, the "best" reward isn't always the most advanced version of the product. It's often the one that offers something only backers can get—whether that's influence, access, or a unique experience. When designing the rewards, consider both the material and immaterial value.

Final Takeaway

Reward-based crowdfunding can offer more than just capital—it's a multi-stage process that can validate ideas, build early communities, and generate public awareness. However, success is far from guaranteed. Many campaigns fail to reach their funding goals, and few products make it to market. To improve the chances, we advised focusing on growing a broad base of backers. Early momentum, clear communication, and compelling, exclusive rewards are key

to standing out in a crowded landscape. Crowdfunding is not just about funding—it's about mobilising the crowd.

References

- [1] Yu, S., & Fleming, L. (2022). Regional crowdfunding and high tech entrepreneurship. *Research Policy*, 51(9), 104348.
- [2] Bessière, V., Stéphanie, E., & Wirtz, P. (2020). Crowdfunding, business angels, and venture capital: an exploratory study of the concept of the funding trajectory. *Venture Capital*, 22(2), 135-160.
- [3] Junge, L. B., Laursen, I. C., & Nielsen, K. R. (2022). Choosing crowdfunding: why do entrepreneurs choose to engage in crowdfunding?. *Technovation*, 111, 102385.
- [4] <https://www.businessresearchinsights.com/market-reports/crowdfunding-market-100103>
- [5] <https://www.kickstarter.com/help/stats?ref%3Dfooter>
- [6] Leland, J. (2024). Kickstarter Data, Global, 2009-2023. Retrieved from: <https://www.icpsr.umich.edu/web/ICPSR/studies/38050>
- [7] Van Teunenbroek, C., Bijvank, J., & Wenzlaff, K. (in print). Prioritising reach over contribution size: On the drivers of market entry after crowdfunding campaign success. *Journal of Alternative Finance*.
- [8] Van Teunenbroek, C., Dalla Chiesa, C., & Hesse, L. (2023). The contribution of crowdfunding for philanthropy: A systematic review and framework of donation and reward crowdfunding. *Journal of Philanthropy and Marketing*, 28(3), e1791.



GenAI Masterclass

A guide to AI-augmented studying in Business Administration

By: Pauline Weritz & Johannes Dahlke

Assistant Professors Faculty of Behavioural, Management and Social Sciences

In recent years, business administration programs have increasingly seen technology as an integral part of management education. While tech-heavy courses are traditionally housed within Computer Science or Business Information Technology at UT, today's business administration students also need digital skills. This includes a technical understanding of technology but should not be limited to that. AI literacy means pairing technical knowledge with critical and reflective thinking about the application domain. The GenAI Masterclass, developed for bachelor's International Business Administration (IBA) and master's Business Administration (BA) students, is one such initiative: an experiment in learning with AI.

Unlike conventional IT courses that emphasize programming, the GenAI Masterclass is about AI literacy and responsible use. Students are not trained to become developers, but to understand, evaluate, and apply generative AI in their academic and professional contexts. Each participant was provided with a license to a generative AI program, ensuring hands-on exploration rather than abstract discussion. This design encouraged students to make AI part of their workflow without turning it into a shortcut. The course was structured into four progressive stages: Foundational knowledge, pre-application setup, active use, post-application reflection:

The GenAI Masterclass placed emphasis on trying, testing, and reflecting. The guiding principle was not "using AI correctly," as the pace of technological progress renders specific procedural tips

(e.g., for prompt engineering) quickly outdated. Instead, it was about learning how to make AI work for yourself in a smart and responsible way. Students worked through a mix of video lectures, tutorials, and self-study exercises. Reflection diaries, augmented by AI itself, captured the evolving experiences of participants, helping them identify strengths, blind spots, and growth opportunities. In this way, the course doubled as both a skill-building exercise and a source of feedback for further development.

Highlights included video demonstrations of how participants combined multiple AI tools in their academic workflows. These demos revealed not just technical skills, but also creativity in tailoring AI support to individual study practices. For us as educators, it offered a glimpse into the future of assessment: moving beyond standard essays toward personalized outputs such as podcasts, videos, and AI-assisted reflections.

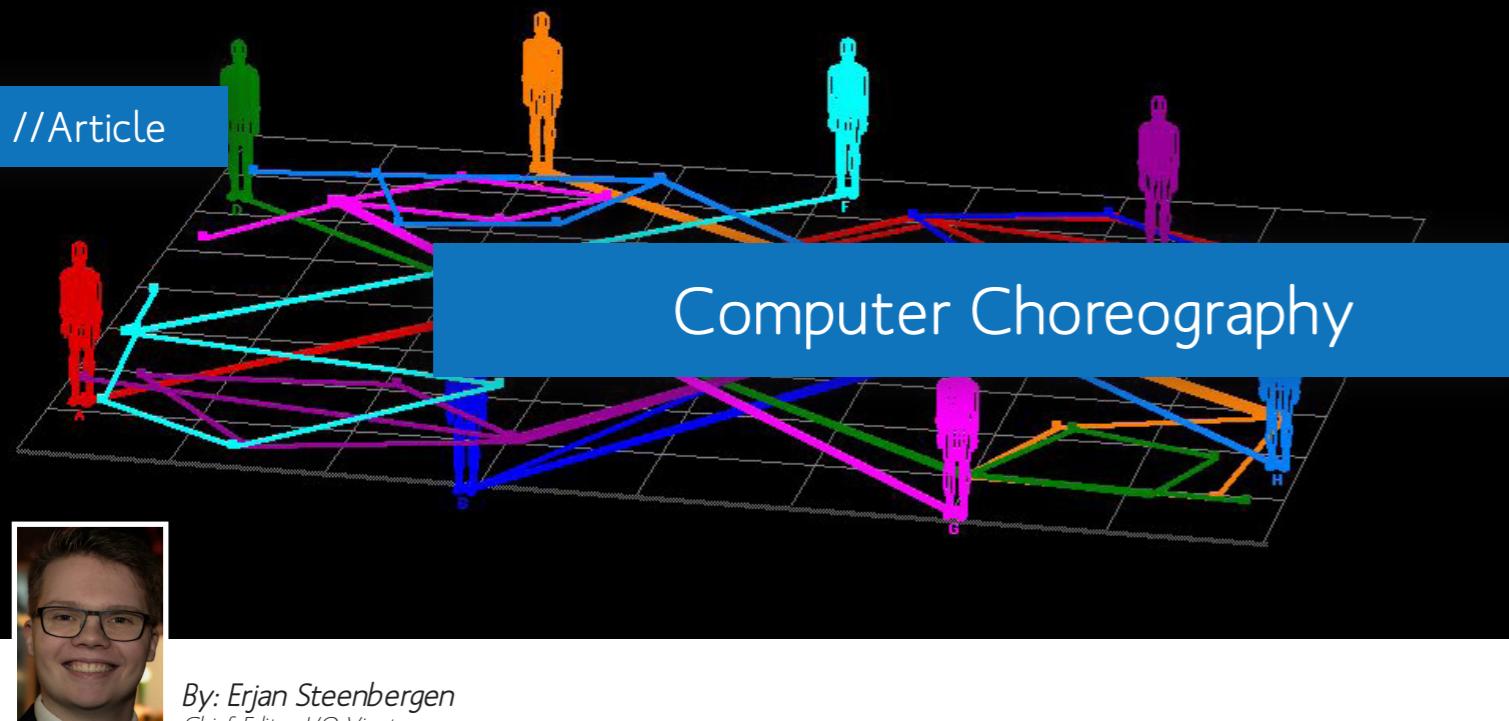
Traditional business education sometimes struggles to balance standardized evaluation with space for personal expression. The GenAI Masterclass showed that generative AI can unlock new forms of assessment. Students experimented with producing multimodal assignments, integrating insights from AI with their own critical reflections. Participants delivered work that bore the imprint of individual experimentation and, in this way, demonstrated how technology can actually enhance, rather than dilute, authenticity in academic work. The success shows that modern education needs to balance learning objectives and assessments that allow space for AI support, while demanding

human expression.

For future managers, consultants, or entrepreneurs, the ability to navigate AI responsibly is becoming as vital as financial literacy. The GenAI Masterclass prepares students to use AI as a sparring partner in idea development, critically evaluate AI outputs instead of accepting them at face value, integrate digital tools into workflows without compromising academic integrity, and reflect on the broader ethical and professional implications of AI adoption. These skills are transferable far beyond university, into consulting presentations, strategic analysis, or any role where digital decision-making is the norm.

Looking Ahead

Funded by WSV, the course has so far been offered as an extracurricular option. Its success suggests room for expansion, potentially embedding AI literacy into the core curriculum of other programs. The Masterclass demonstrates that tech-related education in business administration need not replicate computer science training. Instead, it can focus on building the mindset and skills to work with emerging technologies responsibly and creatively. No hollow promises, no shortcuts, just a space to get inspired, experiment, and reflect. That, perhaps, is the essence of preparing Business Administration students for an AI-augmented future.



By: Erjan Steenbergen
Chief Editor I/O Vivat

"It would be possible to describe everything scientifically, but it would make no sense; it would be description without meaning, as if you described a Beethoven Symphony as a variation of wave pressure."

- Albert Einstein

Science and art are often seen as opposites in society. The coexistence of reason and emotion has brought up debates in both the past and present. This is especially apparent now between AI and art. There is, however, another form of 'art' which has currently not been as tainted by the rise of artificial intelligence as painting and drawing: Dancing.

It might seem strange to even bring it up. How would it even be possible to recreate choreographies with artificial influence? Naturally, as curious our species is, there have already been efforts to use software and AI to help the choreographer come up with interesting dance moves. Nevertheless, the question remains whether computers can create coherent and aesthetically pleasing dance choreography on their own, and if so, how this might affect the dancing industry in the future.

Typically, choreographers develop dances by blending improvisation, notation, and practice with dancers. Some begin with music, crafting movements

that align with the rhythm, while others may focus on a theme or emotion, translating it into physical form. This creative journey often entails a process of trial and error. Choreographers frequently need to physically experiment with sequences alongside dancers, making adjustments to ensure that the choreography appears cohesive. The intricacies of formations and gestures significantly heighten the complexity of choreography.

Challenges emerge as this process can be labor-intensive, time-consuming, and reliant on the creativity of one or a few individuals. At times, choreographers find it difficult to push beyond established movement vocabularies. Additionally, the documentation and sharing of choreography remain problematic.

While systems like Labanotation are available, they tend to be intricate and seldom utilized outside academic circles. In this context, software tools of

fer promising solutions: enabling quick testing of movement concepts, visualizing patterns, and digitally preserving choreographies.

Computer-generated choreography refers to the method of employing algorithms to design dance routines. It is often characterized as using computers for choreographing performances, generating computer animations, analyzing or instructing elements of human movement, demonstrating dance actions, or aiding in the notation of dances. Additionally, it can be utilized in choreographic software for stimulation, facilitating real-time choreography and generative dance, or simulating virtual dancers within the realm of Dance technology.

In the 90s, Merce Cunningham created several choreographies assisted by the software tool Danceforms (then called Lifeforms). Cunningham proclaimed,

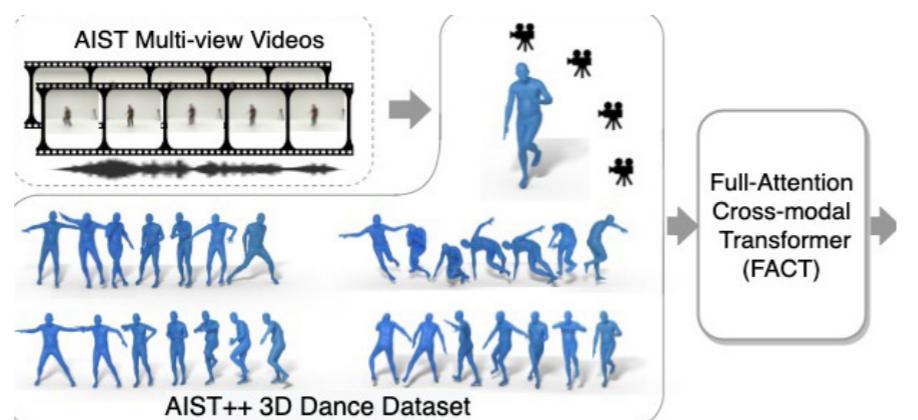


Figure 1: How Google AI Choreographer Works

"It seems clear that electronic technology has given us a new way to look. Dances can be made on computers, pictures can be punched out on them, why not a notation for dance that is immediately visual?", which intuitively seems like a very novel idea that could heavily affect the method of creating dance. Sure, the technological software at that time was unable to give hyperrealistic and fluid movement, but it still got the point across. Given some time and further advances in technology, computer-generated choreography could be a widely used technique.

The choreo itself looks and feels robotic. It has a certain charm to it, however, as the software not only tries to copy human movement, but deliberately integrates a form of movement that looks and feels like a computer. It is therefore quite ironic that in this play, humans replicate the robotic movement, as if we are the ones trying to fool computers that we are one of them.

Circling back to Cunningham's quote, it seems clear that he sees the program more as a tool, used in conjunction with a choreographer, not as a replacement. Danceforms gives dancers a visual representation to help them better understand a play.

The Google AI Choreographer, however, takes a more literal approach to using computers to create choreography. It uses a 3D human dance motion dataset, called AIST++, to artificially generate rhythmically correlated moves

on music. The results are nothing to scoff at, especially compared to previous attempts.

However, is it actually useful (and more importantly, fun) to use these synthetic methods for expressing dance? Does it not defeat the purpose of using creativity to create a choreography by yourself? It is not a tool anymore; it is its own choreographer.

Dance is often valued as a deeply personal act of expression, and creating choreography is part of that artistic journey. The question most people ask is: does delegating the creative part to a machine bypass the essence of dance itself? Others might embrace AI as a collaborative partner though. A tool that sparks new ideas rather than replaces creativity.

Has the pioneering of Cunningham in the 90s actually increased the use of software tools like Danceforms? Currently, the dance industry has not yet found a way to successfully integrate such practices.

One reason computer-generated choreography has not become a staple in the industry is that it does not solve a pressing problem. Unlike lighting control or stage effects, where technology clearly adds efficiency, choreography thrives on human interpretation, collaboration, and intuition. The limitations of machine-generated movement (sometimes rigid, sometimes uncanny) highlight how central the human ele-

ment is to dance.

Furthermore, in professional dance settings, the role of the choreographer is as much about leadership, coaching, and artistic vision as it is about inventing steps. A computer might generate movement patterns, but it cannot guide dancers through the subtleties of performance. This explains why Cunningham and others used tools like Danceforms not as a sketchpad to expand their imagination rather than dictate the final work.

References

- [1] <https://research.google/blog/music-conditioned-3d-dance-generation-with-aist/>
- [2] <https://doi.org/10.1080/01472526.2019.1575661>
- [3] <https://www.computationalcreativity.net/iccc2016/wp-content/uploads/2016/01/Generative-Choreography-using-Deep-Learning.pdf>
- [4] <https://cedardance.com/dance-the-computer-merce-cunningham/>
- [5] https://www.researchgate.net/publication/363529641_The_Computer_A_Choreographer_Aesthetic_Responses_to_Computer-Generated_Dance_Choreography
- [6] https://www.sfu.ca/~tschiphlo/publications/Schiphorst_M.A.Thesis.pdf
- [7] <https://danceinteractive.jacobspillow.org/merce-cunningham-dance-company/trackers/>



Figure 2: Choreography Software LifeForms



From the Chairwoman

"I don't think a board year would fit you"



*By: Jeske van Hijum
Chairwoman Inter-Actief Board & Editor I/O Vivat*

"I don't think a board year would fit you."

Is something that I have heard more than once. Truth be told, I have also heard the opposite often enough, but surprisingly this was always said by the people that would very much benefit from a new person becoming active within the study association. Choosing for a board year calls for introspection, as an objective view is hard to find. People who have done a board year will tell you it is one of the most fun, but also difficult, years of their life. On the other hand, some people who have not done a board year just see a year full of drinking and partying without any real benefits to it. It's safe to say that even in the process of considering doing a board year, there is plenty to consider, and plenty to learn about yourself.

I have always liked to keep myself busy with all types of different things. Quite quickly after starting my studies in Technical Computer Science, I joined the committee market of Inter-Actief to see if I could put myself to use for the association while also having some fun! I joined the Theme Drink Committee and the I/O Vivat in my first year, and quickly decided I wanted to do even more. A part-time board year at Stichting Borrelbeheer Zilverling followed and some more Inter-Actief committees like the RIALLY, the Symposium Committee and the Christmas Committee.



Now, having gotten quite a taste for activism being active, board would be a logical next step in the pipeline, but still some doubts remained. Stopping your studies for a whole year makes for quite the study delay, and it's difficult to be sure about committing to something when you don't yet know how it will turn out. I ended up putting myself past this hurdle, finding the only way to find out if board was actually something for me was to just try it.

So then, despite what I thought I would do for 1,5 years, I became a candidate board member along with 5 other amazing people. We were now recognizable by our beautiful ties for the whole campus as the ones who would (hopefully) be the people running the association.

The time where we used to study was now fully occupied

by preparing ourselves for the year that was to come. Besides getting worked in by our predecessors, we also worked on our policy plan and tried to identify the points on which we want to try and improve the association. Working on the policy plan is the first moment in your candidate board period where you are actively working with your new team to plan out the year you have ahead of you and make it concrete. It is somewhat tough, as everyone going into the discussion has a different view on how things should be done and why, but it is this discussion that brings everyone to new insights, making it one of the most valuable experiences I have had in the period leading up to becoming board.

Next to the working hours, there were of course also the relaxing activities such as an Active Members Week-



end, where we got to connect with the members more by barbecuing for them and organising a round of beer themed hide-and-seek. At the end of our candidate board period we got to join a fun board weekend, where on the first day we got to **REDACTED**. Not to mention, the day after we not only had to **REDACTED**, which was great, but we also did **REDACTED**. Of course, a **REDACTED** could also not miss from a weekend like this. It was definitely one of the most intense, but also fun, weekends I have ever experienced in my life!

After the summer holiday had passed, there it finally was, the day we had all been dreading for weeks: the Change GMM. Would it take 2 hours? Would it take 4? Would it take all night? Would we even get switched in as a board? Luckily, we were done just before midnight, making it a not too long, but also not too short GMM. The feeling of finally being board is kind of like a happy sense of impending doom: you know what difficulties will hit you and you know what challenges (and hangovers) you will face already the day after the GMM, but now you are finally actually in charge! You can finally get to work on the points that you have been preparing to work on for weeks, which is an incredibly exciting feeling.

Now that the first few weeks of being a board have passed, it is nice to look back on these days and reflect on what we have already been able to do for the association. Our officer of internal affairs organized a beautiful committee market where a lot of interest for our committees was generated and we as a

board got to talk to a lot of first- and second-years who have an interest in becoming active members. The New Harvest Drink was organized, which was also a great place to get to know more first-years and talk to them about their experiences of the first few weeks at this university. As the doom from the happy sense of impending doom subsided, I realised that, while it may have been stressful at times, we already succeeded in doing a little bit of usefulness for the association, which was great.

I cannot say that the remark about a board year not being fit for me never got in my head. In fact, it still does sometimes. You don't just breeze through a board year, as some people might think: people expect a lot from you and will ask a lot from you, even when you might not have the time and space for it. This becomes clear even in the first few weeks of being board. But it is at those difficult moments you learn expectation, stress and crisis management, and so many other useful life skills that you just don't get out of your regular life, whether that be your study or your job. And often, even when you have had a hard day, you can look back at it and see that you did keep the association alive yet another day, which also counts for something. I can say with confidence that even just the candidate board period taught me a lot about myself and the way I function under pressure and within a team, which is something I have already gotten to put into use in my first few weeks of being board. I am ecstatic to grow even more in the upcoming year and to do it along with the rest of my amazing board.



Good Idea

If you have any questions for us at the board or ideas for us to improve the association, feel free to contact us at any time at board@inter-actief.net.



By: Ruben Groot Roessink
Editor I/O Vivat

Recently, I took a deeper dive into Linux namespaces as I found that I could solve a CTF (Capture the Flag) challenge using a PID (Process ID) namespace to bypass a PID check in a Linux binary that was part of the CTF. I have done most of my work so far with full virtualization software such as VMWare/Virtual-Box. Containerization software (such as Docker/LXC), which use Linux namespaces, however are interesting concepts that I wanted to learn more about.

In Linux, namespaces are a kernel feature that allow the user to separate/isolate system resources among processes. This means that some processes see only a subset of the mount points, other process IDs, network interfaces, and so on, rather than the full range of them available on the system. This allows us to simulate, for example, a Doc-

ker container only having access to its own isolated process tree. When using `ps aux` within the Docker container, one only sees the processes belonging to that Docker container, whereas other processes must exist on the system that are used by the Operating System running the Docker container.

We will be following the example in Figure 1. We first execute `ps aux` on our full system and pipe the output to the `wc -l` command (short for word count) to count all the lines returned by our command. This shows that we have about 300 processes running on our full system. We then use the `unshare` command to run a program in a modified namespace environment. The following flags were used in our example:

- > `-pid`: Creates a new PID namespace.
- > `-fork`: Tells the command to fork the specific program (`/bin/bash`) to

```
(kali㉿kali)-[~]
$ ps aux | wc -l
300

(kali㉿kali)-[~]
$ sudo unshare --pid --fork --mount-proc /bin/bash
[root@kali ~]
# ps aux | wc -l
4

(root@kali)-[~]
# ps aux
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START    TIME COMMAND
root          1  0.0  0.0  9004  5628 pts/2      S  19:56  0:00 /bin/bash
root          7  0.0  0.0  9104  3572 pts/2      R+ 19:57  0:00 ps aux
 6953  0.0  0.0  21912  7228 pts/1      S+ 19:56  0:00 sudo unshare --pid --fork --mount-proc /bin/bash
 6955  0.0  0.0  21912  2692 pts/2      Ss 19:56  0:00 sudo unshare --pid --fork --mount-proc /bin/bash
 6956  0.0  0.0  5344  1604 pts/2      S  19:56  0:00 unshare --pid --fork --mount-proc /bin/bash
 6957  0.0  0.0  9004  5628 pts/2      S+ 19:56  0:00 /bin/bash
```

Figure 1: Network namespace effects on processes seen

Separation is Hard

Using Linux Namespaces



run as a child process of `unshare` rather than running it directly.

> `-mount-proc`: Just before running the program, it mounts the `proc` filesystem at `/proc`, starting a new mount namespace in the process. This ensures that no information can be retrieved from other processes outside of the namespace (as most tools will interact with `/proc` to retrieve information about processes).

We run the `unshare` command with `/bin/bash` as the program to fork (starting a new root shell in the process, due to the fact that we need to run `unshare` with root permissions by default). When we now list the number of processes we see that the number of processes that our current shell can observe has dramatically decreased. Listing the processes gives us two processes: PID 1 with `/bin/bash` and PID 7 as our `ps aux` process.

Listing the same processes outside of the newly created namespace, we can see three processes (alongside many others) that are used by `sudo/unshare` to start the new namespace and `/bin/bash` that is the PID in our new namespace (See Figure 1).

The following eight namespaces exist (since kernel version 5.6):

- > Mount (mnt) - Control of mount points

```
(kali㉿kali)-[~]
$ ls -lah /proc/self/ns
total 0
dr-x---x-- 2 kali kali 0 Jul 29 20:32 .
dr-xr-xr-x 9 kali kali 0 Jul 29 20:32 ..
lrwxrwxrwx 1 kali kali 0 Jul 29 20:32 cgroup → 'cgroup:[4026531835]'
lrwxrwxrwx 1 kali kali 0 Jul 29 20:32 ipc → 'ipc:[4026531839]'
lrwxrwxrwx 1 kali kali 0 Jul 29 20:32 mnt → 'mnt:[4026531841]'
lrwxrwxrwx 1 kali kali 0 Jul 29 20:32 net → 'net:[4026531840]'
lrwxrwxrwx 1 kali kali 0 Jul 29 20:32 pid → 'pid:[4026531836]'
lrwxrwxrwx 1 kali kali 0 Jul 29 20:32 pid_for_children → 'pid:[4026531836]'
lrwxrwxrwx 1 kali kali 0 Jul 29 20:32 time → 'time:[4026531834]'
lrwxrwxrwx 1 kali kali 0 Jul 29 20:32 time_for_children → 'time:[4026531834]'
lrwxrwxrwx 1 kali kali 0 Jul 29 20:32 user → 'user:[4026531837]'
lrwxrwxrwx 1 kali kali 0 Jul 29 20:32 uts → 'uts:[4026531838]'
```

Figure 2

> Process ID (pid) - Independent set of process IDs (starting with PID 1 again)

> Network (net) - Virtualization of the network stack

> Inter-process Communication (ipc) - Isolation of processes to communicate with other processes in other namespaces (See Figure 2). While only grasping the surface of namespaces in this article (both limited by the actual number of words allowed and my understanding of namespaces yet), this hopefully shines some light on the internals of the Linux operating system that we make use of every day and many might have never seen/cared about before.

> UTS (Unix Time-Sharing) - Isolate host and domain names

> User ID (user) - Allows us to isolate privileges and user identification (this allows us to become root in a Docker container, while this user does not have root permissions outside of the container)

> Control group (cgroup) - Hides identities of the control group of which the process is a member. Control groups are used to allocate, limit and monitor the resources used by processes

> Time - Isolate system times to allow different processes to work with different times

By default, a single instance of a Linux operating system starts with a single instance of each of the namespaces. Depending on the need, new namespaces

can be started. The User ID and the PID namespaces are nested as children into the namespaces that created them.

We can list the namespaces of processes (that we have access to) using the `/proc/{PID}/ns/` directory, which contains soft links to specific namespaces (See Figure 2). While only grasping the surface of namespaces in this article (both limited by the actual number of words allowed and my understanding of namespaces yet), this hopefully shines some light on the internals of the Linux operating system that we make use of every day and many might have never seen/cared about before.

I would like to end this article with the example that made me look into this (and that you can follow if you like). The CTF challenge as put forward had a check for the PID of the process. If the PID matched (a) certain value(s), a flag was printed. If not, the program would kill itself after telling the user to try again. I have a very simple example for this explanation, the actual challenge had to be reverse-engineered (e.g. using Ghidra).

The following C-code in Figure 3 takes the following steps:

- > It imports some C header files (containing definitions of data types and

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

int main() {
    const pid_t TARGET_PID = 7;
    pid_t current_pid = getpid();
    if (current_pid == TARGET_PID) {
        printf("Success, here is a flag: CTF{IO_Vivat_2025_123456}");
    } else {
        printf("Try again!");
    }
    return 0;
}
```

Figure 3: PID check C program

allowing the use of `getpid()`;

> It starts a main function (which is the entry point when running the final executable);

> It then sets a `pid_t` constant to match against. In this case, the challenge is to ensure that the PID of the process running the application is 7;

> We then retrieve the PID of the current process using the `getpid()` syscall;

> Finally, we compare the two values in an if-statement. If the if-statement is true, our flag is printed, otherwise the user is told to 'Try again!'

Of course, in this particular case it is much easier to get the flag in another way (using the `strings` command or a decompiler for example). In a CTF, the flag would be encoded/encrypted/dynamically built. For the sake of simplicity, this was omitted from the test program.

We execute the following command to start a new namespace, starting our PID counter at 1: `sudo unshare --pid --fork --mount-proc /bin/bash`

During testing, I found that several processes are run after our first command, ensuring that we have used the first 4 PIDs after our initial command. We can increase our PID by one by using `bash &` (this starts a new bash shell in the background). This means that if we do this twice, we end up at PID 6. We can then run our program (`./pid`) to run it as PID 7, meaning that our flag is printed.

Last Note

This challenge can also be solved without using namespaces. One solution is to overwrite the `/proc/sys/kernel/ns_last_pid` value, which contains the last PID assigned in our current namespace (using `echo {PID} > {file}`). But this was non-trivial during testing. I could get the PID within a certain range, but not at the correct value every time. This would have been easier, if the program had checked if the PID was within a certain range. Another solution would have been to hook the `getpid()` function using the `LD_PRELOAD` variable and ensure that this would always return 7.



By: Koen Braham, Kevin Alberts, Tom Grooters, & Josje van 't Padje
ENIAC Board

With a new year comes a new board. ENIAC would like to introduce the new board for 2025. The new board is composed of the following members:

Hi! My name is Koen Braham (31), currently on the ENIAC board as chairman. I graduated back in 2018 for my master Computer Science following the Wireless and Sensor Systems specialization. Unfortunately, this track is no longer offered as the number of students were too few. The specialization offered courses on sensor systems and data

networks. Next to that it gave flexibility to pick and match interesting courses from cyber security and electro engineering. For me this created an interesting playing field with theoretical and practical application of this knowledge.

Hey everyone, my name is Kevin Alberts (31), and I am the current treasurer of ENIAC. I finalized my bachelor in Technical Computer Science in 2021, and have been working full-time since then. Way back in 2015-2016, I was part of the IAPC Foundation board, also handling the treasurer role, and I'm still involved with their Board of Advisors today.

Since my graduation I try to keep up with my friends and fellow alumni via ENIACs activities. Especially the larger activities such as visit to theme parks

are a highlight for me. Being on the board I enjoy helping students on the well-known speed-dates to find a great location for their graduation.

During my studies I joined a variety



Figure 1: ENIAC activity (family activity) at Apenheul

of committees at Inter-Actief, among which the SysCom, WWW-committee and SocCie. After finishing my bachelors, I have been enjoying organizing and visiting the ENIAC activities, to meet old friends again and talk with other alumni I don't get to see often. I also still enjoy helping out with Inter-Actief's technical committees. It's a good way to keep my own IT knowledge up-to-speed by trying and learning new things, and it benefits the association as a plus!

Hey hey! My name is Tom Grooters (29), and my current role in the ENIAC board is the secretary. I finished my Master Computer Science with the Software Technology track in 2023.

I originally started in 2016 with Mechanical Engineering but I quickly switched over to the Computer Science Bachelor in early 2017.

During this time I have been very active at multiple associations. From doing a board year at Inter-Actief as a treasurer, to being in the Cash Audit Committee For SV Dimensie (psychology). And although I am still active at some places, like the SysCom of Inter-Actief and Board of Advisors at IAPC, I wanted to do something more with my free time and joined the ENIAC Board starting this year.

Since graduating I have been hard at work at my current job but still aim to make sufficient time to keep up the con-

nections I have made with some good friends during my study. I have also joined many of the ENIAC activities, starting with the speed-dates as a student, and now I am enjoying organizing the activities for others to join.

I'm Josje van 't Padje (31), currently a General Board Member of ENIAC. I started my bachelor's in Technical Computer Science (back when it was still called Technische Informatica) in 2013.

During my studies, I was active within Inter-Actief, serving as the Officer of Educational Affairs during the 2015-2016 board year. I also joined many committees, both before and after my board year—something I really enjoyed and look back on fondly.

After finishing my bachelor's, I continued with the Computer Science master's, following the Data Science track, which I completed in 2021.

Not long after graduating, I realized how much I missed the sense of community and involvement. So when I was asked to join the ENIAC board, I didn't hesitate for long. Being part of ENIAC allows me to stay connected with fellow alumni and contribute to keeping the Inter-Actief spirit alive even after graduation. I really enjoy helping to strengthen the bond between alumni and the association, and it's great to stay in touch with familiar faces, and meet new ones, through our events and initiatives.



Figure 2: ENIAC computer, the namesake of the association.

ENIAC

ENIAC is the alumni association for the studies Business & IT and Technical Computer Science and corresponding masters. Additionally, ENIAC represents the interests of the previous study Telematics. ENIAC was founded to stimulate contact between alumni as well as between alumni and current students. In order to do so, ENIAC organizes several activities (see below) that are both open to members of ENIAC and Inter-Actief.

Activities

21 or 28 November - Drink - Enschede (Campus)

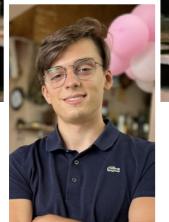
13 December - Annual GMM

16 January - New Year's Drink (Utrecht)



Dutch Education System

Why the Netherlands is unique in high school education



By: Filip Karkalasev
Editor I/O Vivat

The Netherlands is often praised for its innovative approach to education. Nowhere is this more evident than in its voortgezet onderwijs (secondary education), which stands out as one of the most unique high school systems in the world. Dutch secondary education differentiates early, making students pursue an education that fits their academic strengths and career prospects, unlike many countries where all students follow the same high school curriculum. Secondary school in The Netherlands usually starts around age 12.

The tracks are:

- > vmbo (Voorbereidend Middelbaar Beroepsonderwijs): 4-years of secondary school preparing students for the MBO (vocational education).
- > havo (Hoger Algemeen Voortgezet Onderwijs): 5-years of secondary school leading to higher HBO (universities of applied sciences).
- > vwo (Voorbereidend Wetenschappelijk Onderwijs): 6-years of secondary school preparing students for university.

What makes the Dutch system particu-

larly unique is its flexibility. Many secondary schools have combined tracks in the first year of secondary school (vmbo-tl/havo and havo/vwo), making it easier for students to transfer between levels at an early stage. However, also later on, a student who, for example, excels in havo may move up to vwo.

Flexibility continues later on in a student's career with profielen (a profile according to the subjects that a student chooses). Students choose their profile and follow it in the bovenbouw (the upper grades, which are years 3&4 for vmbo, years 4&5 for havo and years 3-6 for vwo). Each profile has a set of compulsory subjects, plus electives. On havo



and vwo the profiles are:

- > Natuur & Techniek (Nature & Technology)
- > Natuur & Gezondheid (Nature & Health)
- > Economie & Maatschappij (Economics & Society)
- > Cultuur & Maatschappij (Culture & Society)

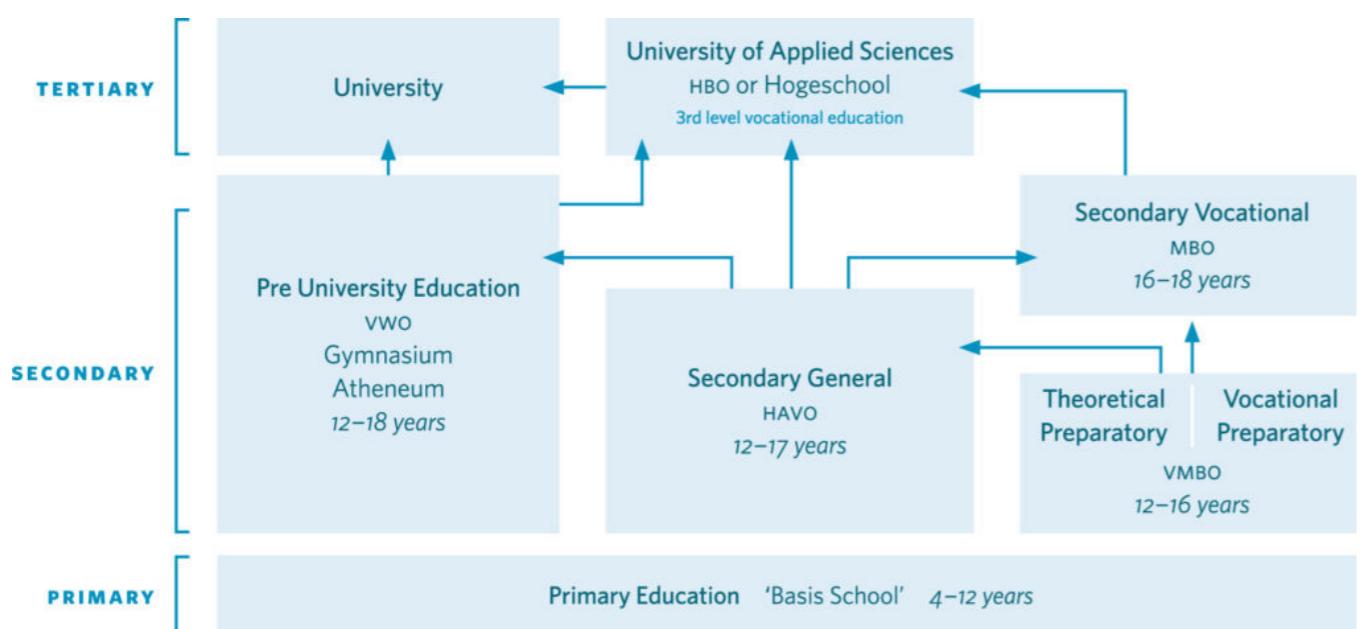
You can follow a mixed profile if you, for example, choose N&T with Biology, which is a common choice, and is aptly called NT/NG. Vmbo has different sublevels (vmbo-tl, vmbo-bb, vmbo-kb and vmbo-gl), with vmbo-tl having 3 of the profiles similar to havo and vwo and the fourth one being agriculture instead of culture & society, and the other vmbo sublevels have more specialized profiles, including care & welfare and business & commerce.

In addition, Dutch schools have an international and globalized attitude. Many Dutch secondary schools offer tweetalig onderwijs or tto (bilingual education), with all subjects (except physical education and foreign languages) being taught in English in the onderbouw (lower grades), offering students the possibility to significantly enhance their English proficiency. Full bilingual programs are less widespread in other European countries. This reflects the Netherlands' international

sis, and the customizable profiles and other adaptability additions, such as flex-hours, add to the Dutch culture of independence and autonomy. Together, these elements make Dutch secondary education distinctive and equip students not just for exams, but for autonomous participation in an increasingly interconnected world.

References

- [1] <https://www.erasmusplus.nl/sites/default/files/2024-02/HE%20system%20and%20scheme%20Dutch%20education%20system%20feb%202024.pdf>
- [2] <https://www.slo.nl/international/the-dutch-education-system/>
- [3] https://english.onderwijsinspectie.nl/binaries/onderwijsinspectie_eng/documenten/publications/2024/06/05/dutch-education-system---background-information/Background+informatiion+Dutch+education+system.pdf





By: *Erjan Steenbergen*
Member Hyper-Actief & Chief Editor I/O Vivat

Sports. Even *Inter-Actief* does it. You might think that we, as computer science students, are not well-known for our sporting prowess, but that is not true! Every year, we participate in the local stAf tournament (we even made it to the finals this year!) and the Batavierenrace. *Hyper-Actief* helps in this by organising the Batavierenrace for the members of *Inter-Actief*. That is not all! We also organise many sporting activities during the year to keep your body fit!

There are many sports clubs on campus. So many that it might be difficult to choose one. We aim to assist you in your choice by organising workshops for *Inter-Actief* members from various associations, especially sports that you would not directly think of. Here are

some of the workshops that we organised in the past:

Spikeball, sometimes called roundnet, is an energetic game played with two teams of two around a small circular net on the ground. One team serves the ball onto the net so it bounces out toward the other team, and from there, both sides try to return it using up to three touches. Once the point begins, there are no sides or boundaries, so players run all the way around the net trying to set up good angles and hit the ball where the opponents can't reach it. A rally ends when a team can't get the ball back onto the net, and because the court is essentially a 360-degree circle, the game tends to be fast, reactive, and a bit chaotic in a fun way.

Floorball is an indoor sport that looks

a bit like a lighter, faster version of hockey. Players use plastic sticks and a hollow plastic ball, and each team has five field players plus a goalie. The game is played on a small rink with boards, and the focus is on quick movement, clean passing, and good stick handling rather than heavy physical contact. Rules limit high sticks and body checks, which makes the sport accessible to a wide range of players. Because the equipment is light and the pace is high, the game flows quickly and rewards good teamwork.

Of course, we not only organise workshops, but also other fun sporting activities, such as Crazy Kangaroo, Bunker Dodgeball, and Beach Volleyball! We always try to organise around 8 activities per year (including the Batavierenrace), so there is always an activity that floats your boat! There are some activities already in the pipeline, so be sure to check the activities on the *Inter-Actief* site every so often!



Hyper-Actief

Check out the committee here:
<https://www.inter-actief.utwente.nl/committees/217/hyper-actief/>

If you would like to organise these activities with us, be sure to hit us up at hyper-actief@inter-actief.net