



AI×MUSIC

Workshop

**Recommenders and Intelligent Tools in
Music Creation: Why, Why Not, and How?**

Speakers

**Christine
Bauer**

Johannes Kepler
University Linz



**Peter
Knees**

TU Wien



**Richard
Vogl**

TU Wien



**Hansi
Raber**

Hansi Raber



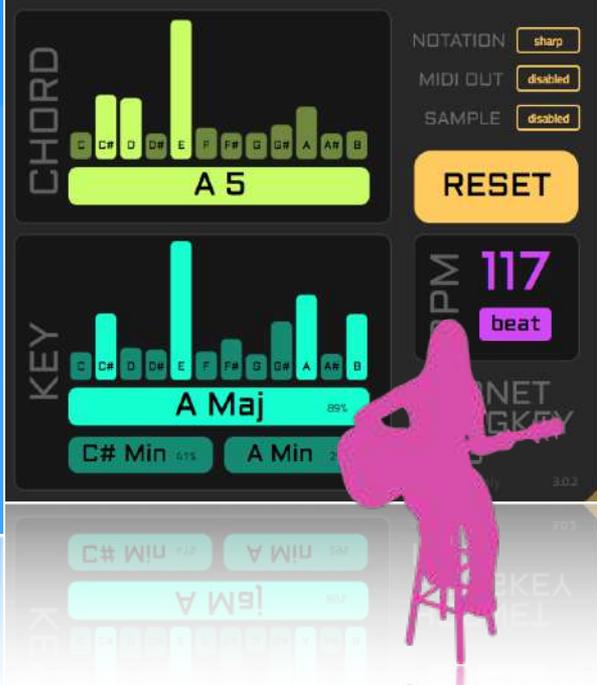
Artificial Intelligence and Music



Recommendations for listeners

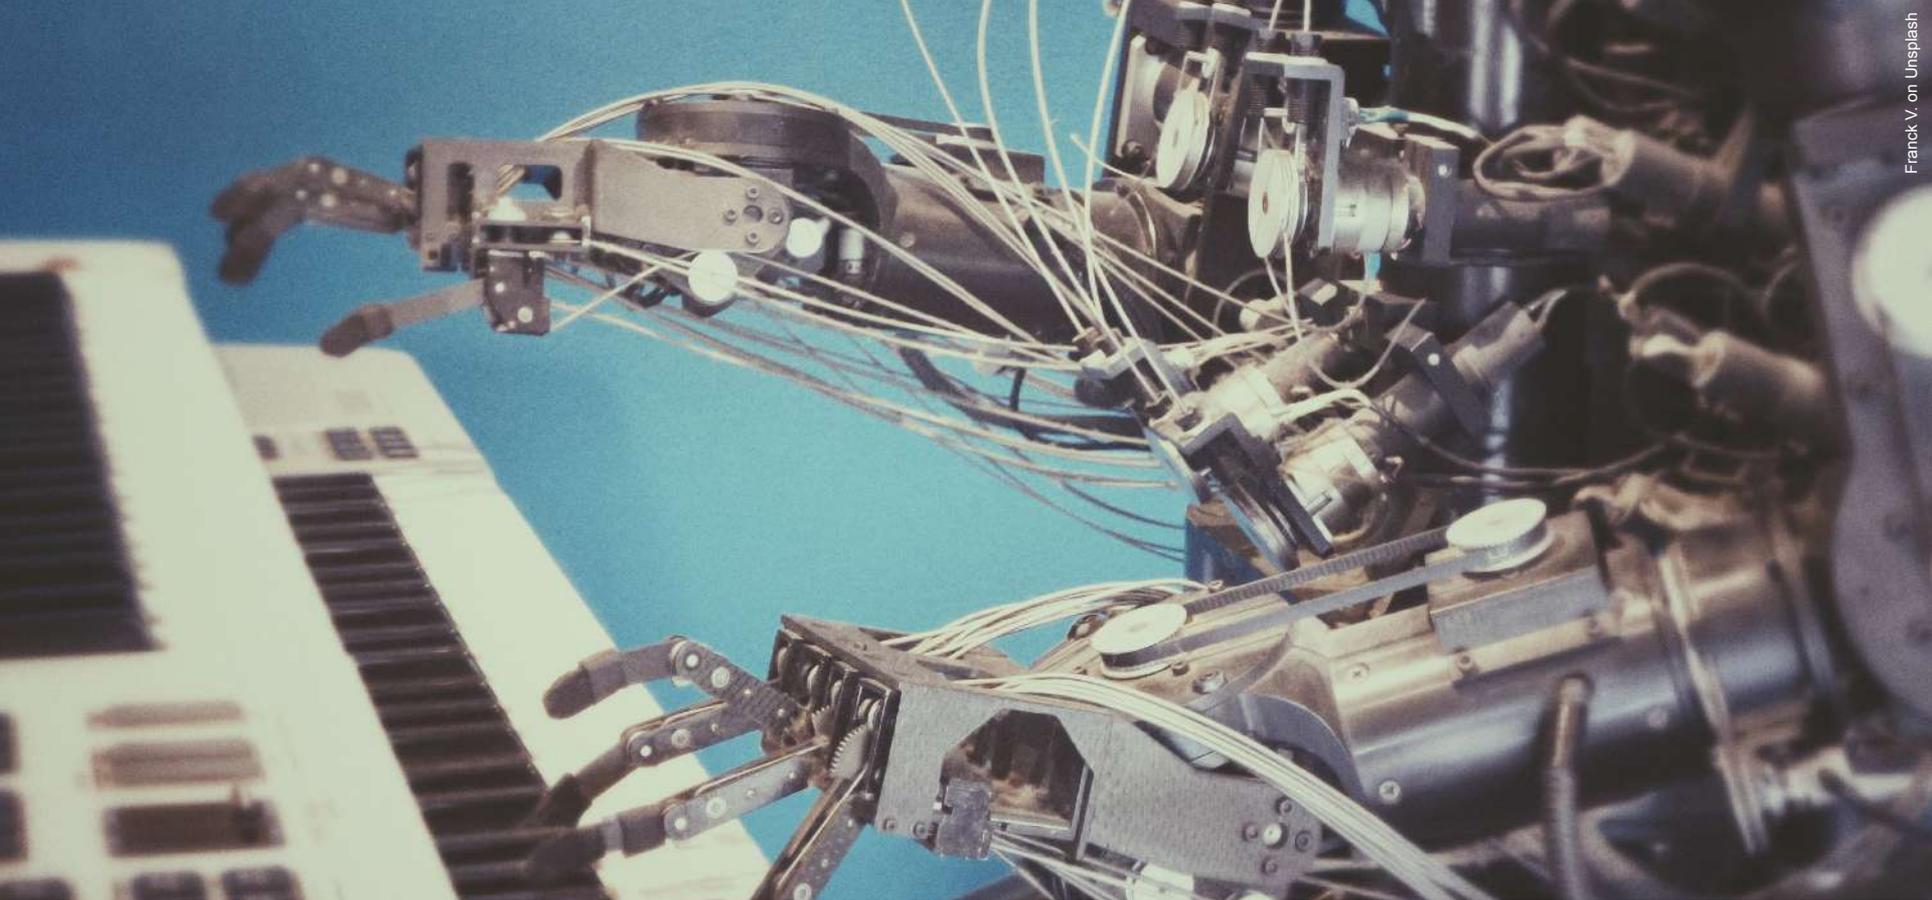


Song identification



Chord and key identification

Typical associations



AI as music creator – the future?



Quantization & Slicing



Melody assistant



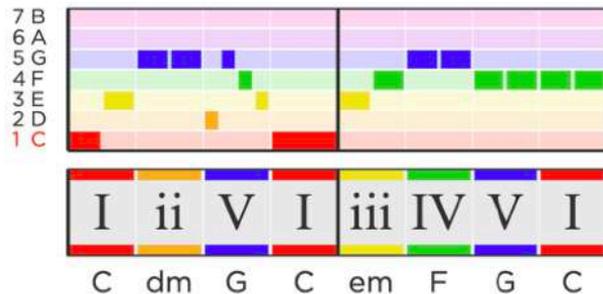
Automatic drumming



Autotune



Automatic mixing



Chord suggestions

Intelligent tools in music creation



**Peter
Knees**



Recommenders and intelligent tools for creative music making



**Christine
Bauer**



Reflection and open questions

Recommenders for Music Makers



Choice overload

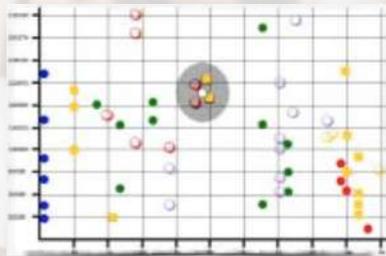
Finding the right sound remains a central challenge



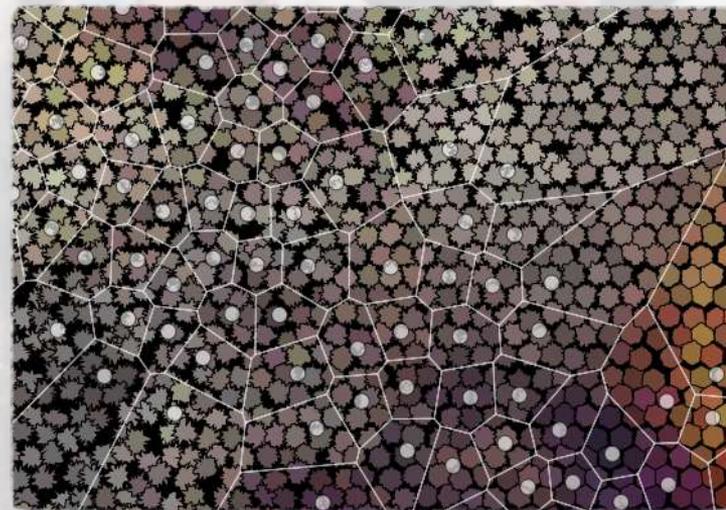
“We usually have to browse really huge libraries [...] that most of the time are not really well organized.” (TOK003)

“[I have] like, two hundred gigabytes of [samples]. I try to keep some kind of organization.” (TOK006)





Sonic browser
(Fernström and
Brazil, ICAD 2001)



Texture browser
(Grill and Flexer, ICMC 2012)

Audio Quilt: snare, synth
(Fried et al., NIME 2014)



Drum sample browser
(Pampalk et al., DAFx 2004)



Sample Browsing Interfaces

User-centric approach

- ▶ participatory workshops, semi-structured interviews
- ▶ conversations with international up-and-coming musicians





Recommenders are seen critical in creative work

***“I am happy for it to make suggestions,
especially if I can ignore them” (TOK007)***

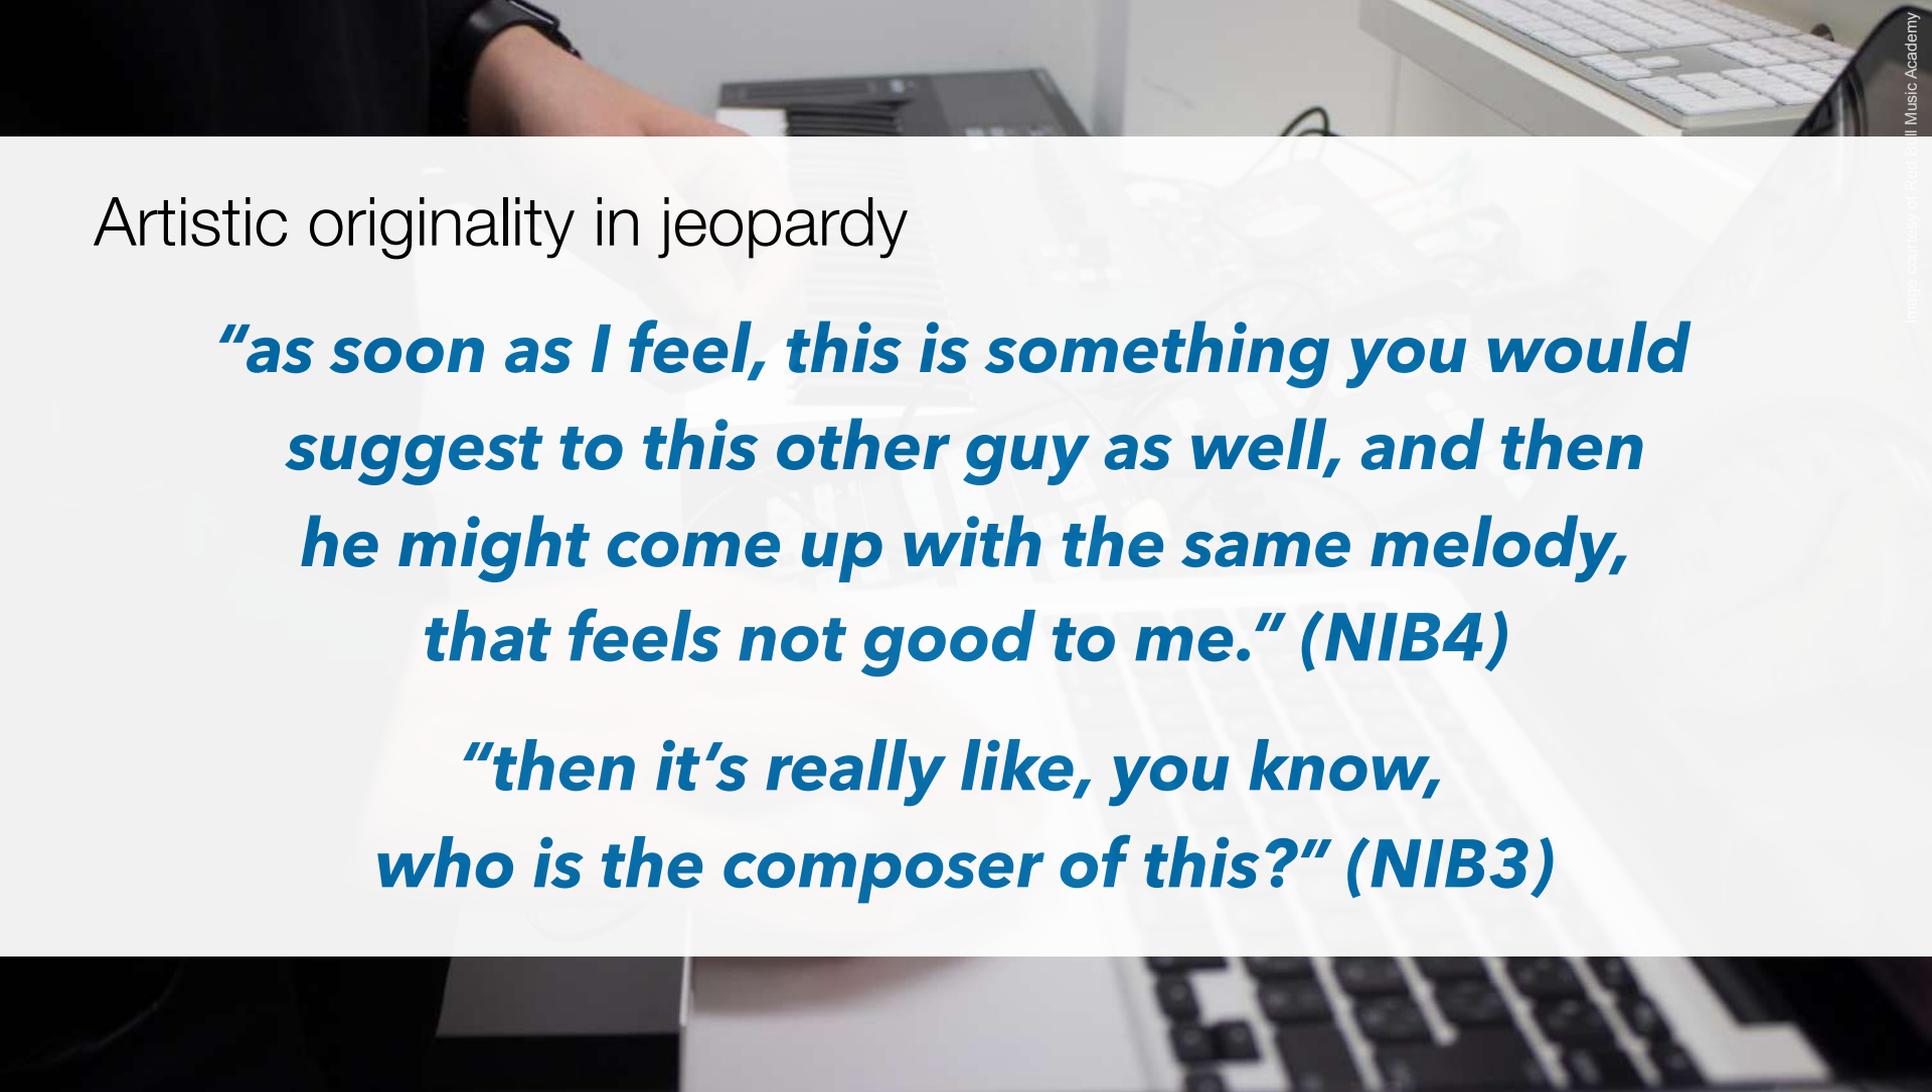




Who is in charge?

***“as long as it is not saying
do this and do that.” (TOK009)***

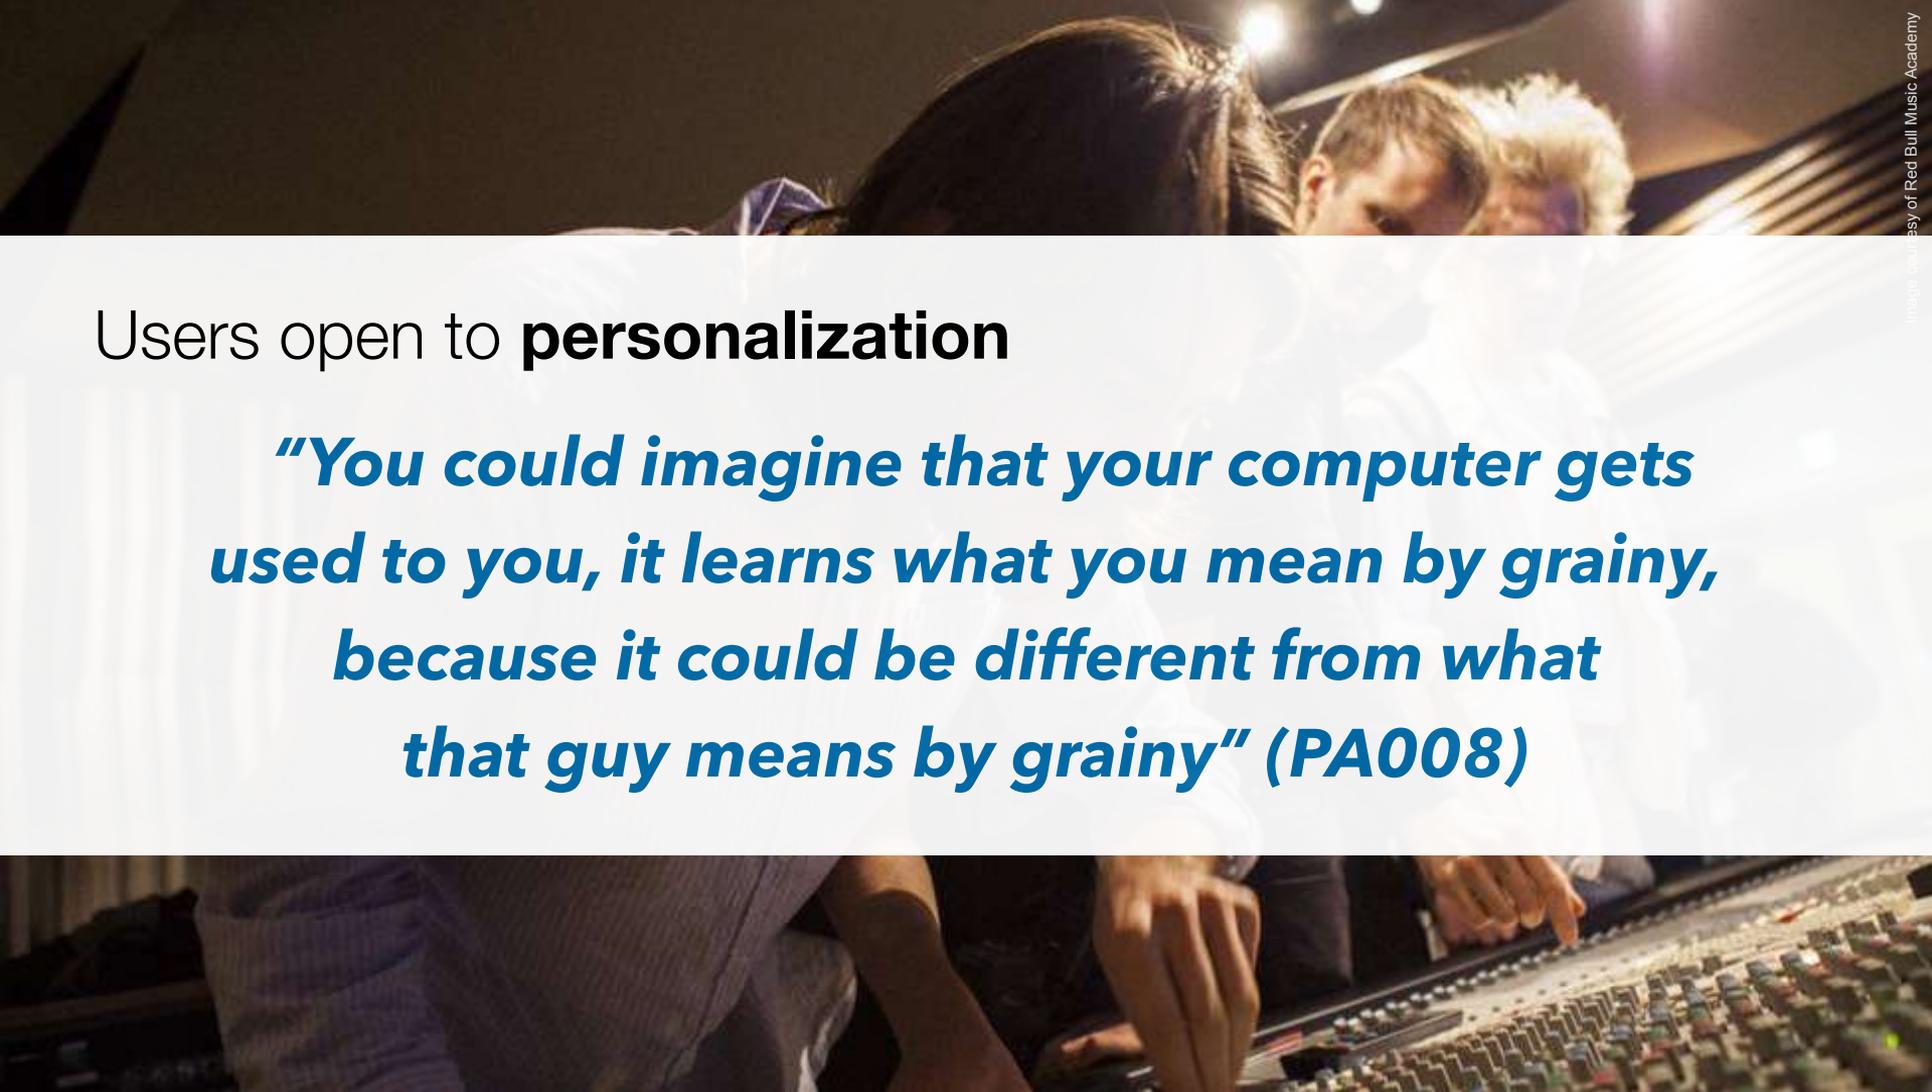




Artistic originality in jeopardy

“as soon as I feel, this is something you would suggest to this other guy as well, and then he might come up with the same melody, that feels not good to me.” (NIB4)

“then it’s really like, you know, who is the composer of this?” (NIB3)

A photograph of a recording studio. In the foreground, a person's hands are visible, adjusting knobs on a large, complex mixing console. In the background, several people are looking down, presumably at a computer screen, in a dimly lit room with some stage lights visible. The overall atmosphere is professional and focused.

Users open to **personalization**

“You could imagine that your computer gets used to you, it learns what you mean by grainy, because it could be different from what that guy means by grainy” (PA008)

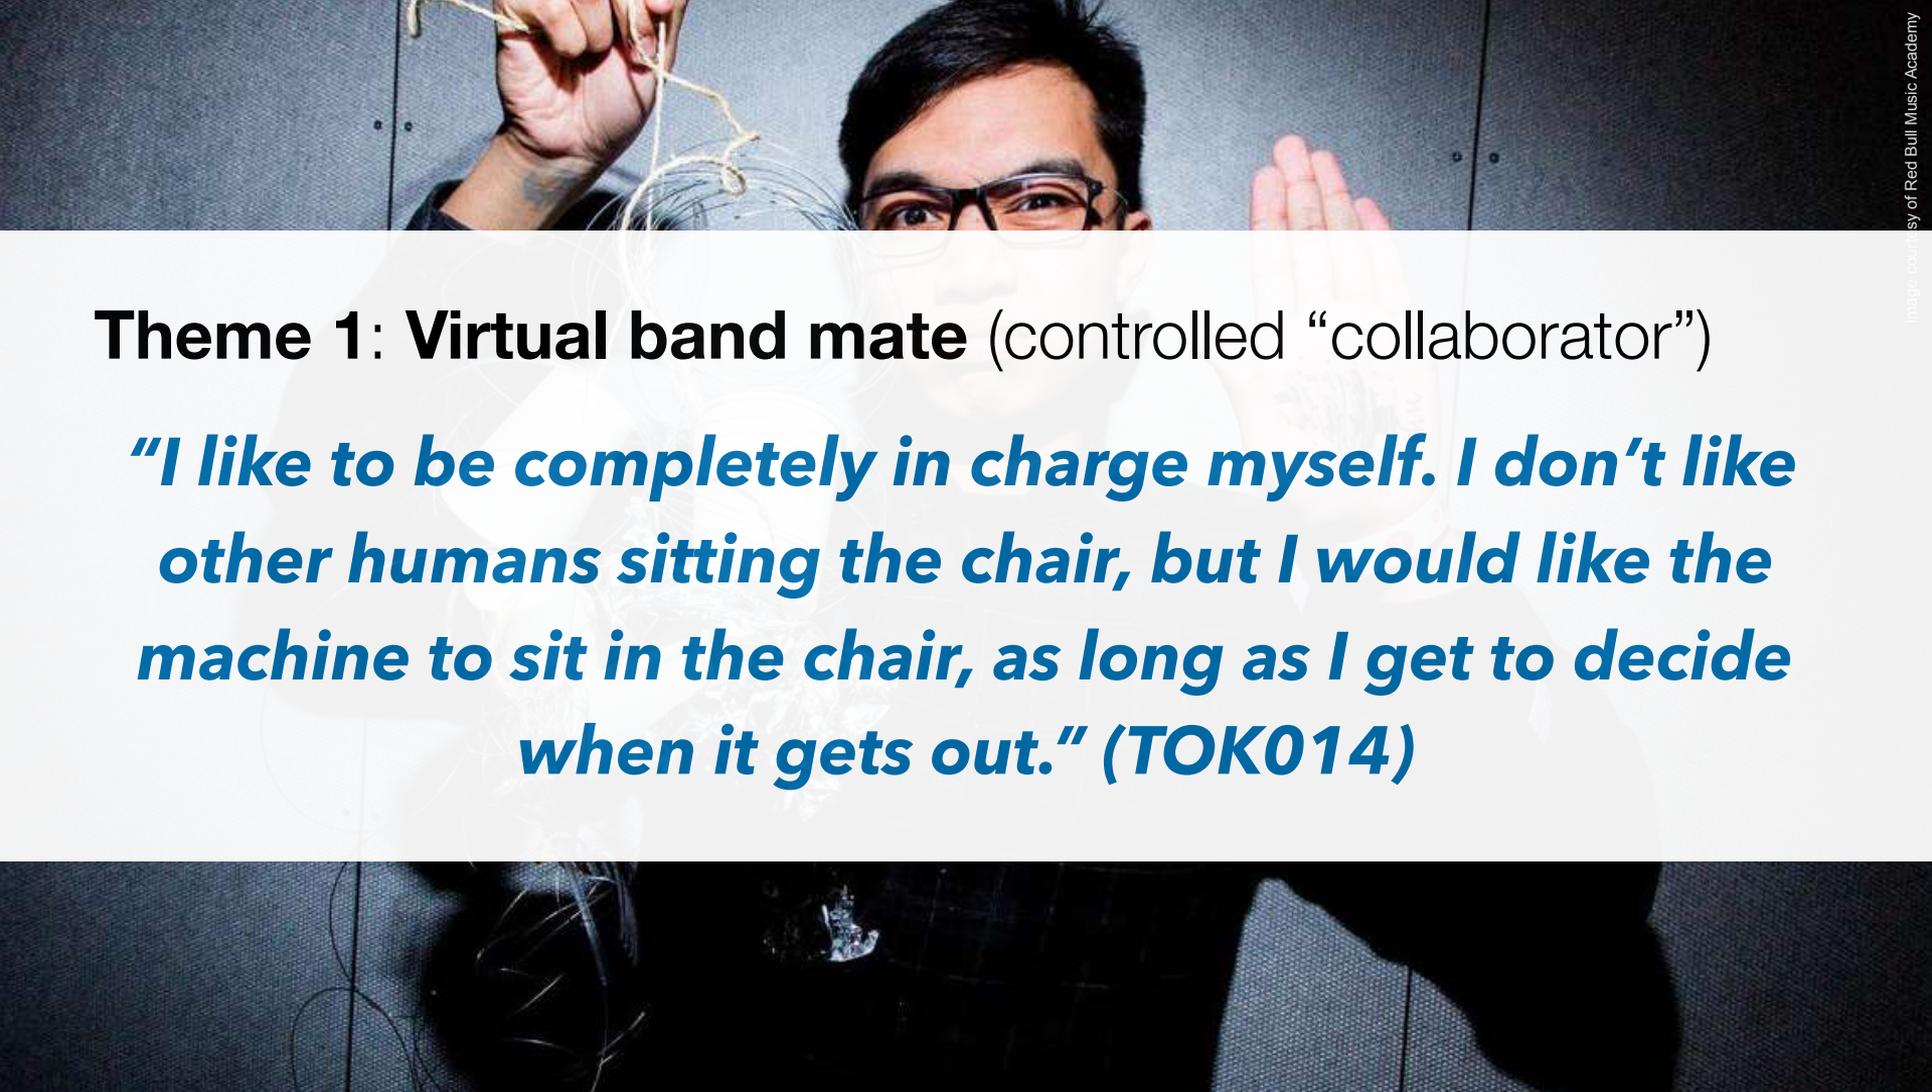


Imitation is not the goal, **opposition** is the challenge

“I’d like it to do the opposite actually, because the point is to get a possibility, I mean I can already make it sound like me, it’s easy.” (TOK001)



“Make it complex in a way that I appreciate, like I would be more interested in something that made me sound like the opposite of me, but within the boundaries of what I like, because that’s useful. Cause I can’t do that on my own, it’s like having a bandmate basically.” (TOK007)



Theme 1: Virtual band mate (controlled “collaborator”)

“I like to be completely in charge myself. I don’t like other humans sitting the chair, but I would like the machine to sit in the chair, as long as I get to decide when it gets out.” (TOK014)

Theme 2: Exploring non-similarity (“the other/strange”)

“So if I set it to 100% precise I want it to find exactly what I am searching for and probably I will not find anything, but maybe if I instruct him for 15% and I input a beat or a musical phrase and it searches my samples for that. That could be interesting.” (TOK003)



The “Other” in Creative Work

- ▶ no interest in imitating existing ideas and “filter bubbles”
- ▶ challenge and question expectations and past behavior



Opposite goals when recommending for creative work

Defamiliarization

Predictability & Explainability

Change of Context

Context Preservation

Opposition

Imitation

Obstruction

Automation



courtesy of Red Bull Music Academy

Takeaways

- ▶ experts need recommenders mostly for **inspiration**
- ▶ a useful recommender needs to be a **collaborator**



Technology Demo: AI Drummers

Why AI Drummers

- As an **inspirational tool**
- Increase **productivity**
- **Use cases:**
 - Music production
 - Live performances
- **Challenges:**
 - Many degrees of freedom
 - Genre dependent
 - No well defined measure of quality
 - Original, meaningful, but not random patterns!



Reactable ROTOR

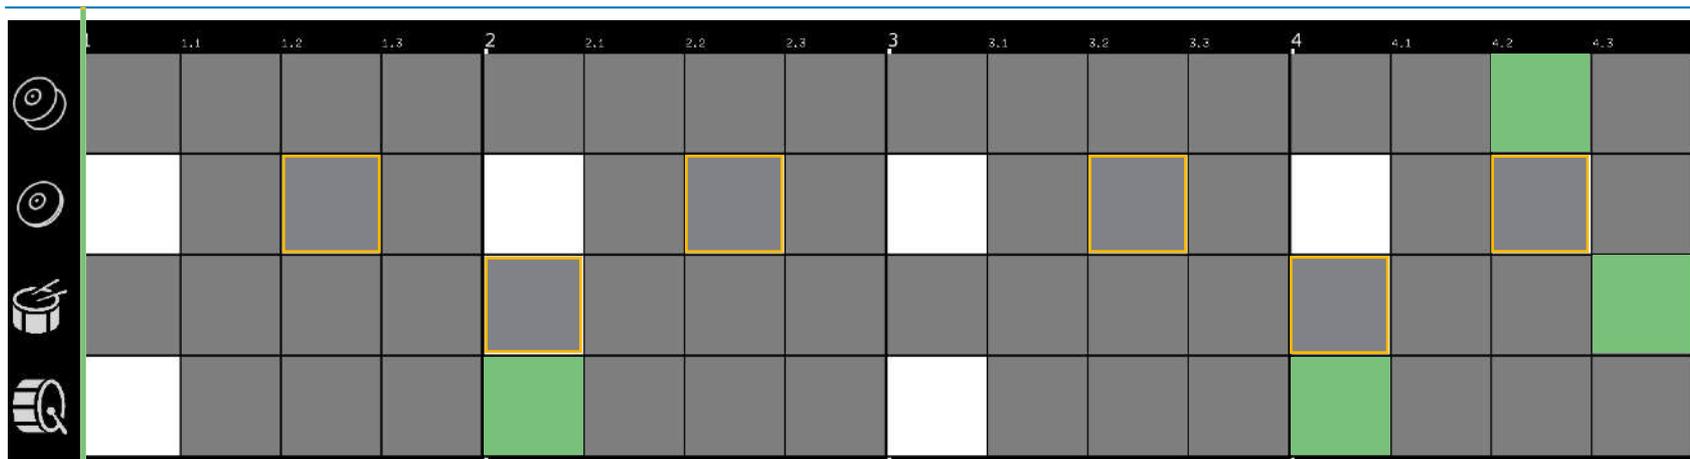


NI Maschine

Examples for AI Drummers

- Drum Pattern Variation
Generate variations of a seed pattern
- Parametric Drum Pattern Generation
Control properties of drum pattern which should be created

Drum Pattern Variation

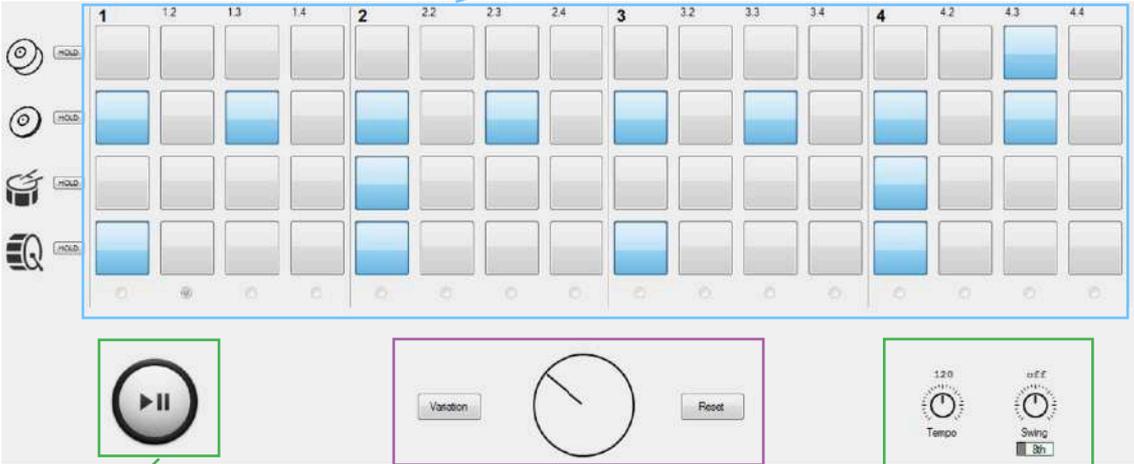


- Create modifications of a given seed pattern
- Maintain characteristic of the drum pattern (the beat)
- **Add details** to increase intensity
- **Remove notes** to make it more simple

This is a step-sequencer...

User Interface

step sequencer grid:
pattern editor and visualization



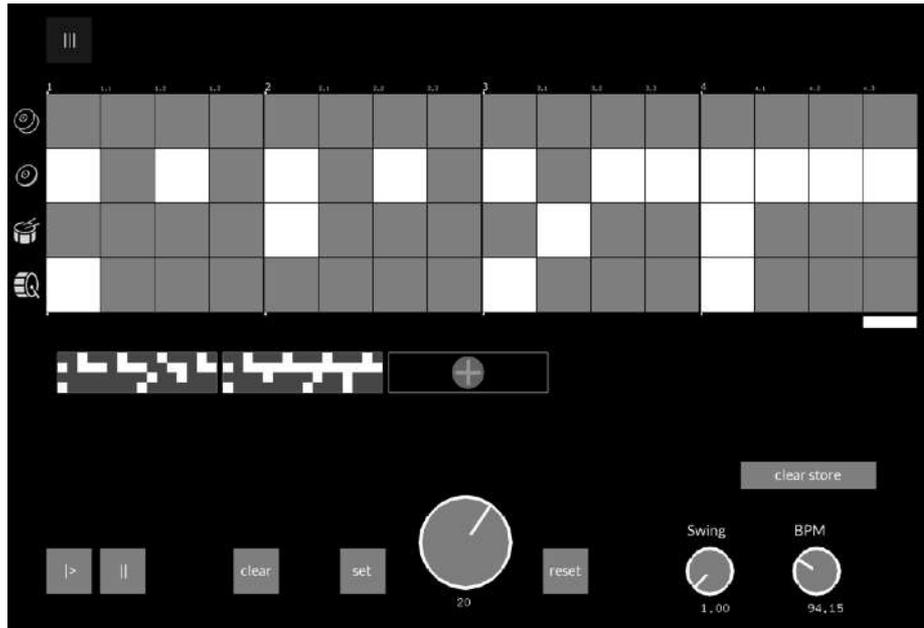
playback controls

generator controls:
pattern browsing dial

tempo and swing controls (playback)



Touch UI



Pattern Variation Method

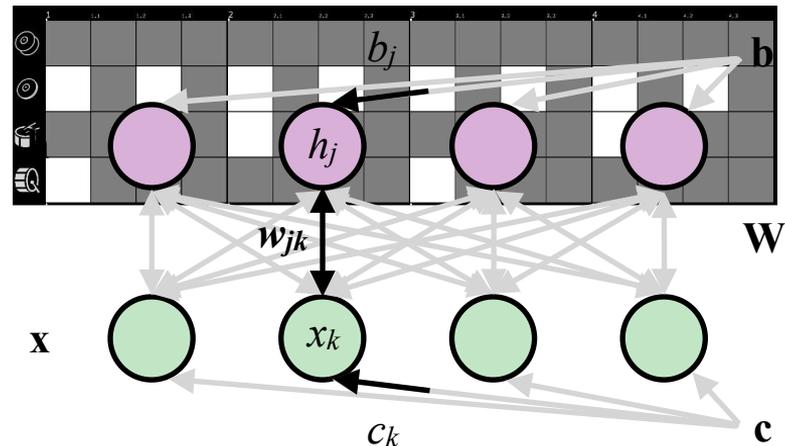
- Focus on **electronic dance music** (EDM)

- **Step sequencer interface**

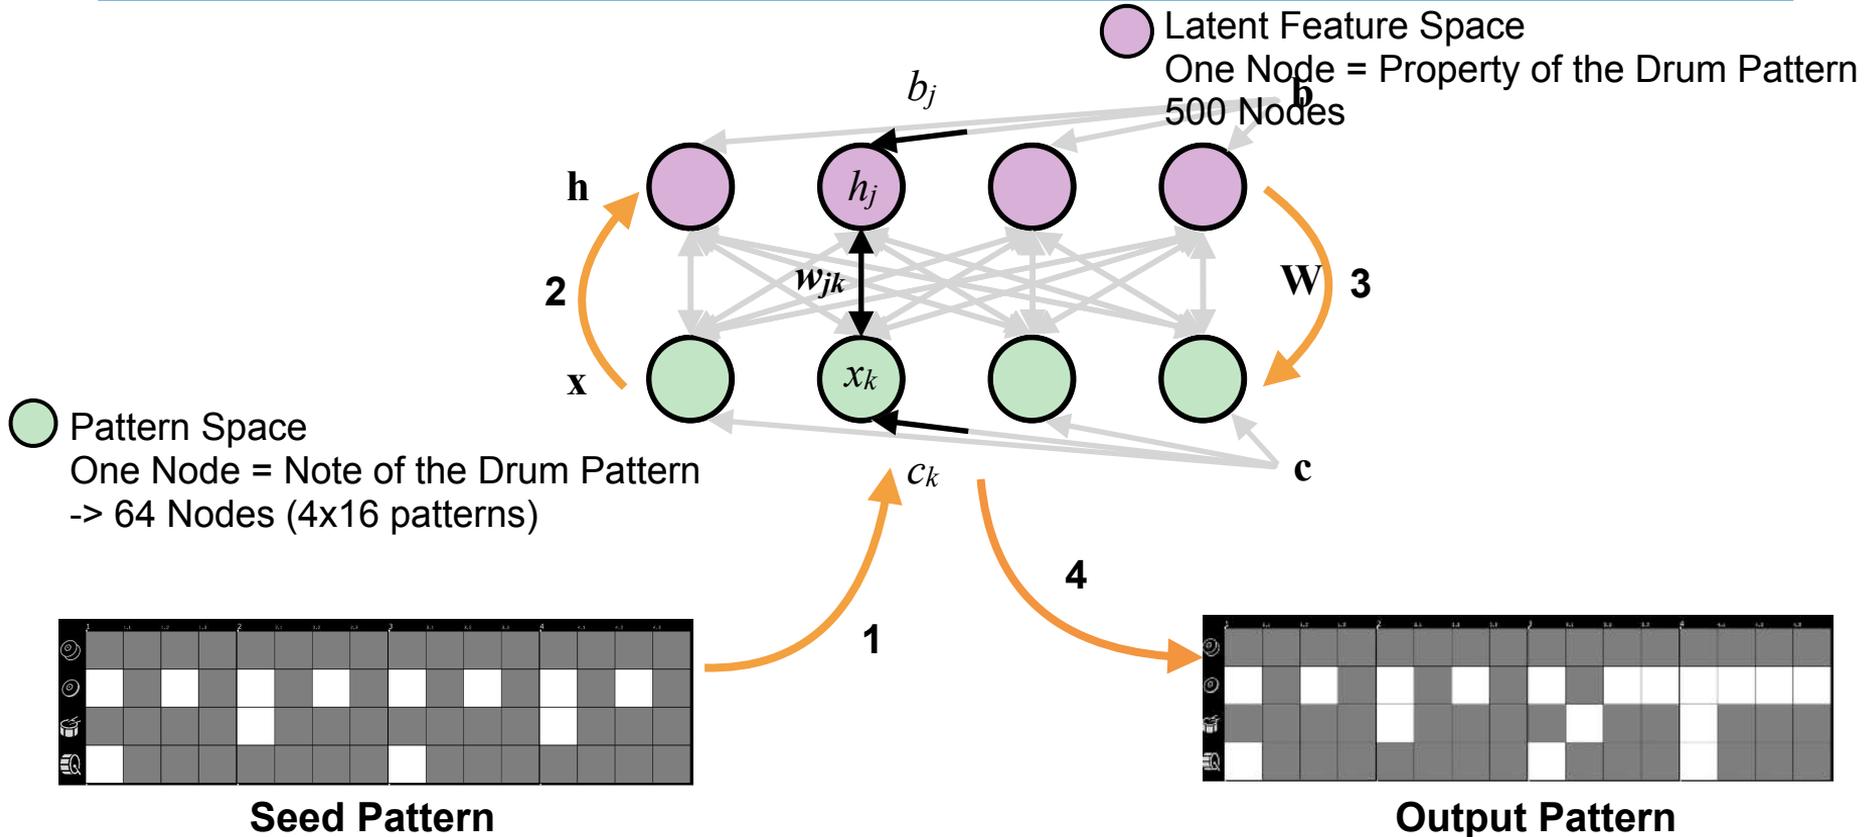
- 4/4 time signature
- 16th note resolution, 4 instruments
- Fixed pattern grid size

- **Stochastic generative model**

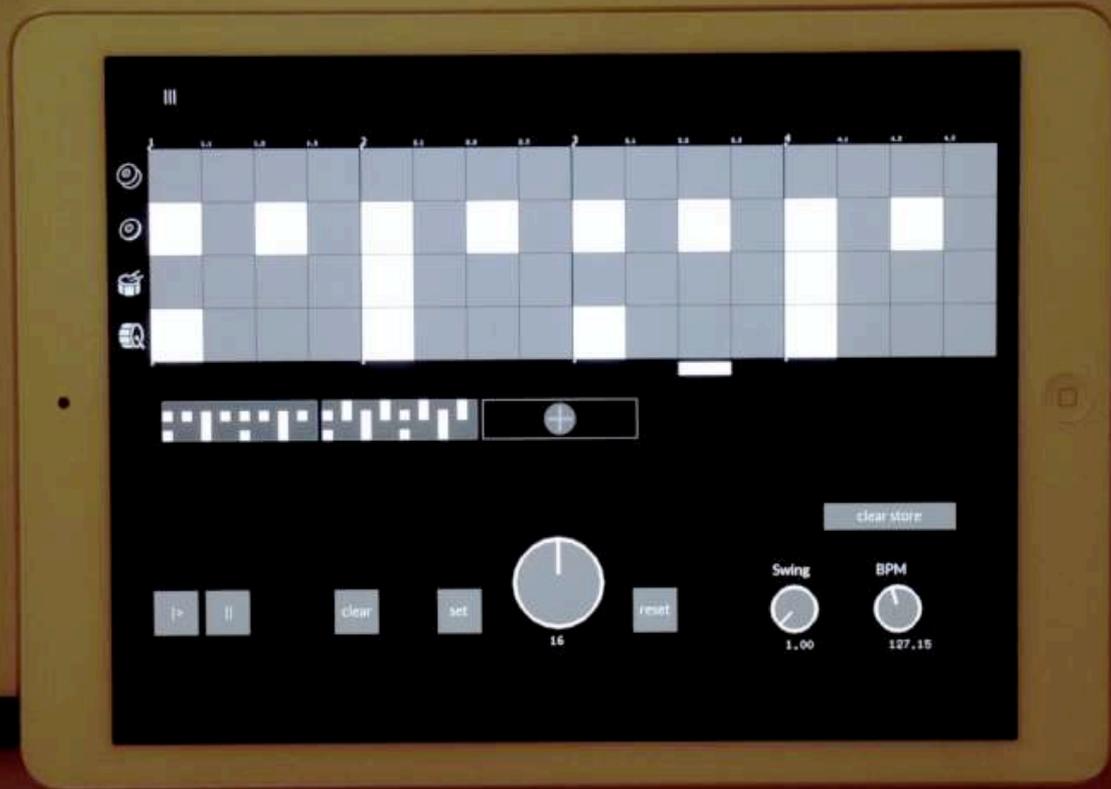
- Sampling of **restricted Boltzmann machine** (RBM)
- Trained on EDM **drum loop library** (NI Maschine)



Pattern Variation Method



Demo



Parametric Drum Pattern Generation

- Create drum patterns given certain properties:
 - Genre
 - Complexity
 - Loudness
- Usually this is done using labeled pattern libraries
 - e.g.: Logic Drummer
 - Often perceived to be unoriginal
- Let's use a stochastic generative model!

User Interface

The image shows a user interface for a music production application. At the top is a large grid labeled "step sequencer grid" with 4 columns and 16 rows. Below the grid is a "bar position indicator" showing a sequence of notes. At the bottom, there are several control panels: "playback controls" with play and stop buttons; "generator controls" with "clear" and "new" buttons and a dial; a "loudness/complexity - pad" with a square pad; and "tempo and swing controls (playback)" with two dials for "Swing" (set to 1.00) and "BPM" (set to 120.00).

step sequencer grid

bar position indicator

playback controls

generator controls

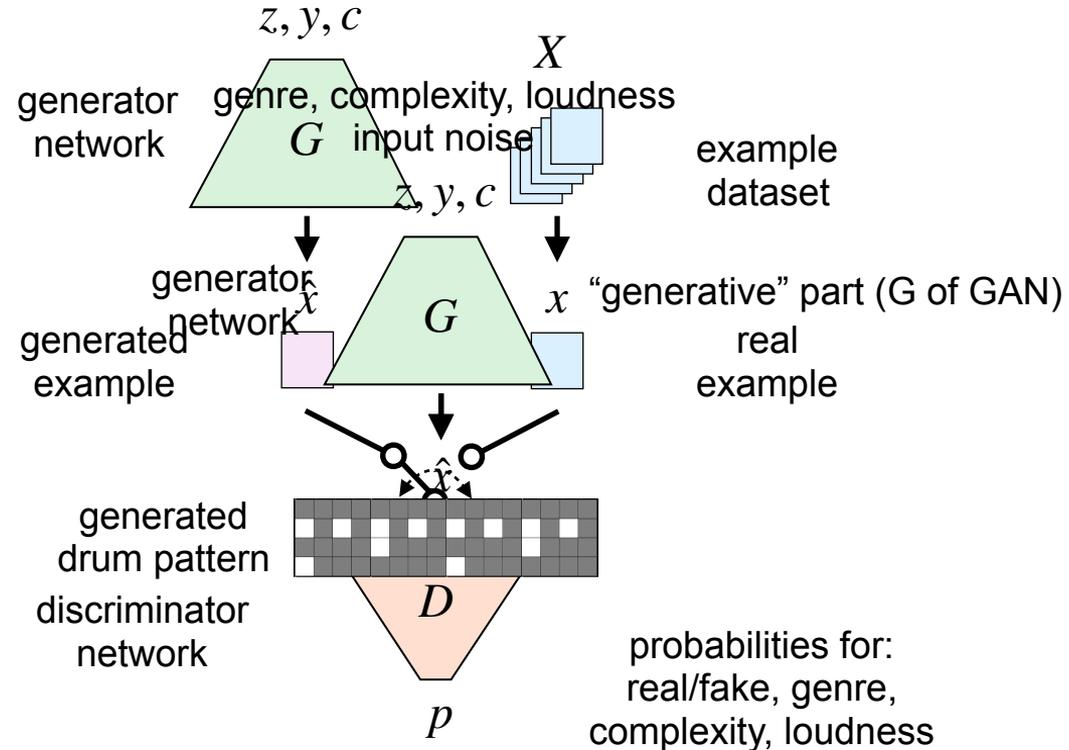
loudness/complexity - pad

tempo and swing controls (playback)

Generation Method

- Train on dataset with **multiple musical genres**
- **Step sequencer interface**
 - 4/4 time signature
 - 32nd note resolution, 8 instruments
- **Generative model**
 - Generative Adversarial Neural Network (GAN)

Pattern Generation Method



Demo

The interface features a piano roll at the top with a grid of notes. The x-axis is labeled with time values from 1.1 to 4.7. The y-axis is labeled with notes C, D, E, F, G, A, B. A vertical menu on the left contains icons for various instruments and effects. Below the piano roll is a mixer with four channels, each with a volume knob and a solo button. A horizontal slider is positioned below the mixer. At the bottom, there are several control elements: a play button, a stop button, a 'clear' button, two circular meters (one labeled '12' and the other '0'), a central square control panel with 'loud', 'soft', 'simple', and 'complex' settings, and two BPM meters labeled 'Swing' (2.89) and 'BPM' (120.55). A 'store' button is located above the BPM meters.

Technology Demo: ClinkyDinky



ClinkyDinky

Audio Flow



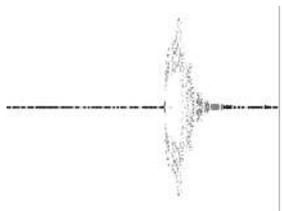
On- and Offset
detection



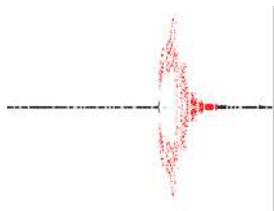
Sample
cropping




Sample
classification



audio stream

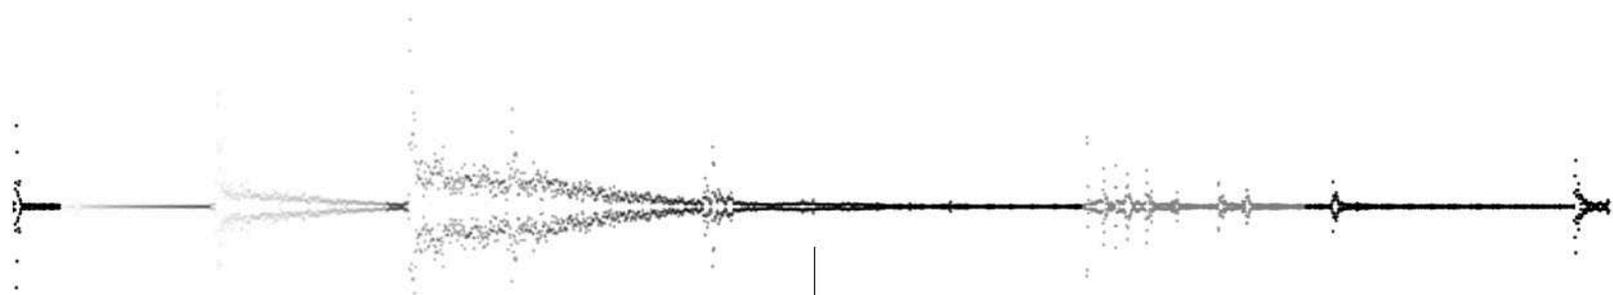


sample



“click”

Wiedergabe aktivieren (Leertaste)		Aufnahmelautstärke				Aufnahme-Archiv (Leertaste)					
Mikrofon aktivieren (r)		Wiedergabelautstärke				Zufällig Geräusche wählen					
Alles löschen (c)		Tempo									
play	rec	clear	mic sensitivity 1.00		volume 1.00		bpm 120		archive	rand	export

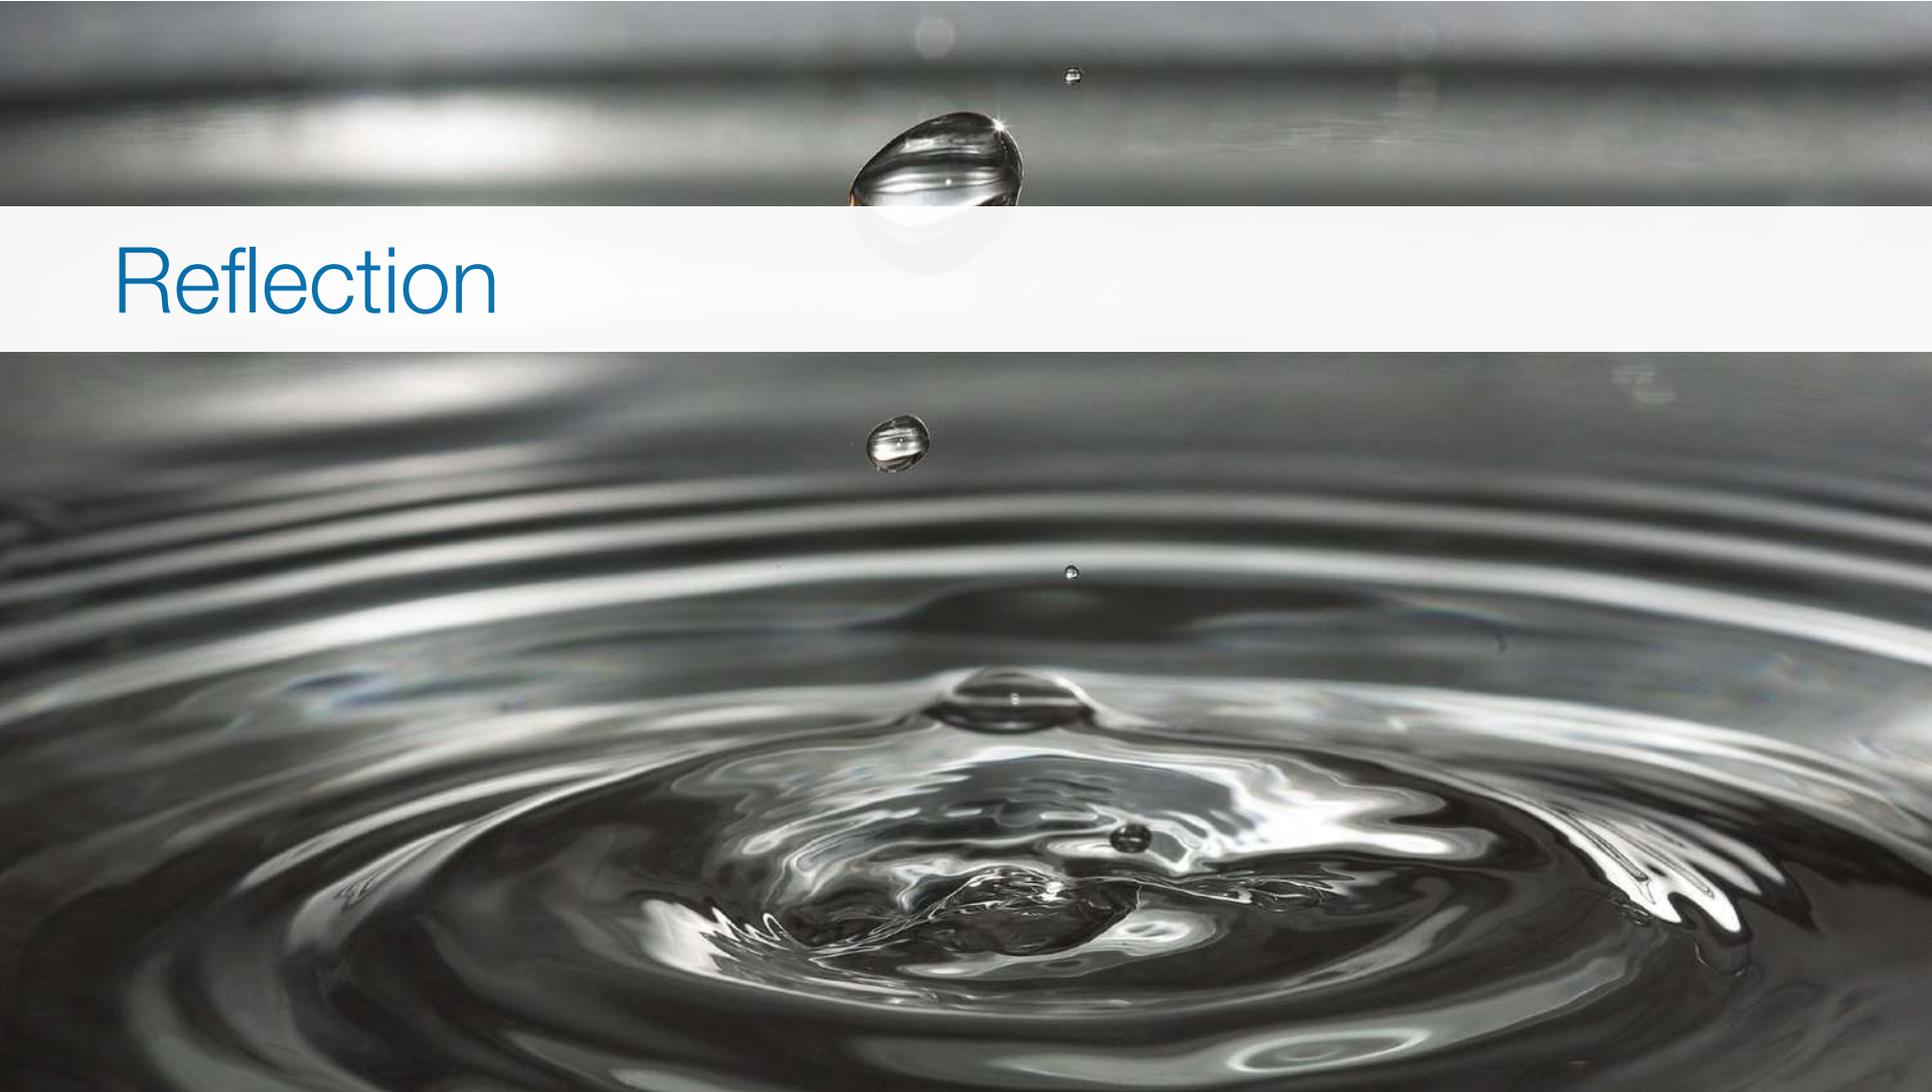


1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
rand	rec														

Nächste Aufnahme hier ablegen
 Mit zufälligen Sample aus dem Archiv belegen
 Step ein/ausschalten

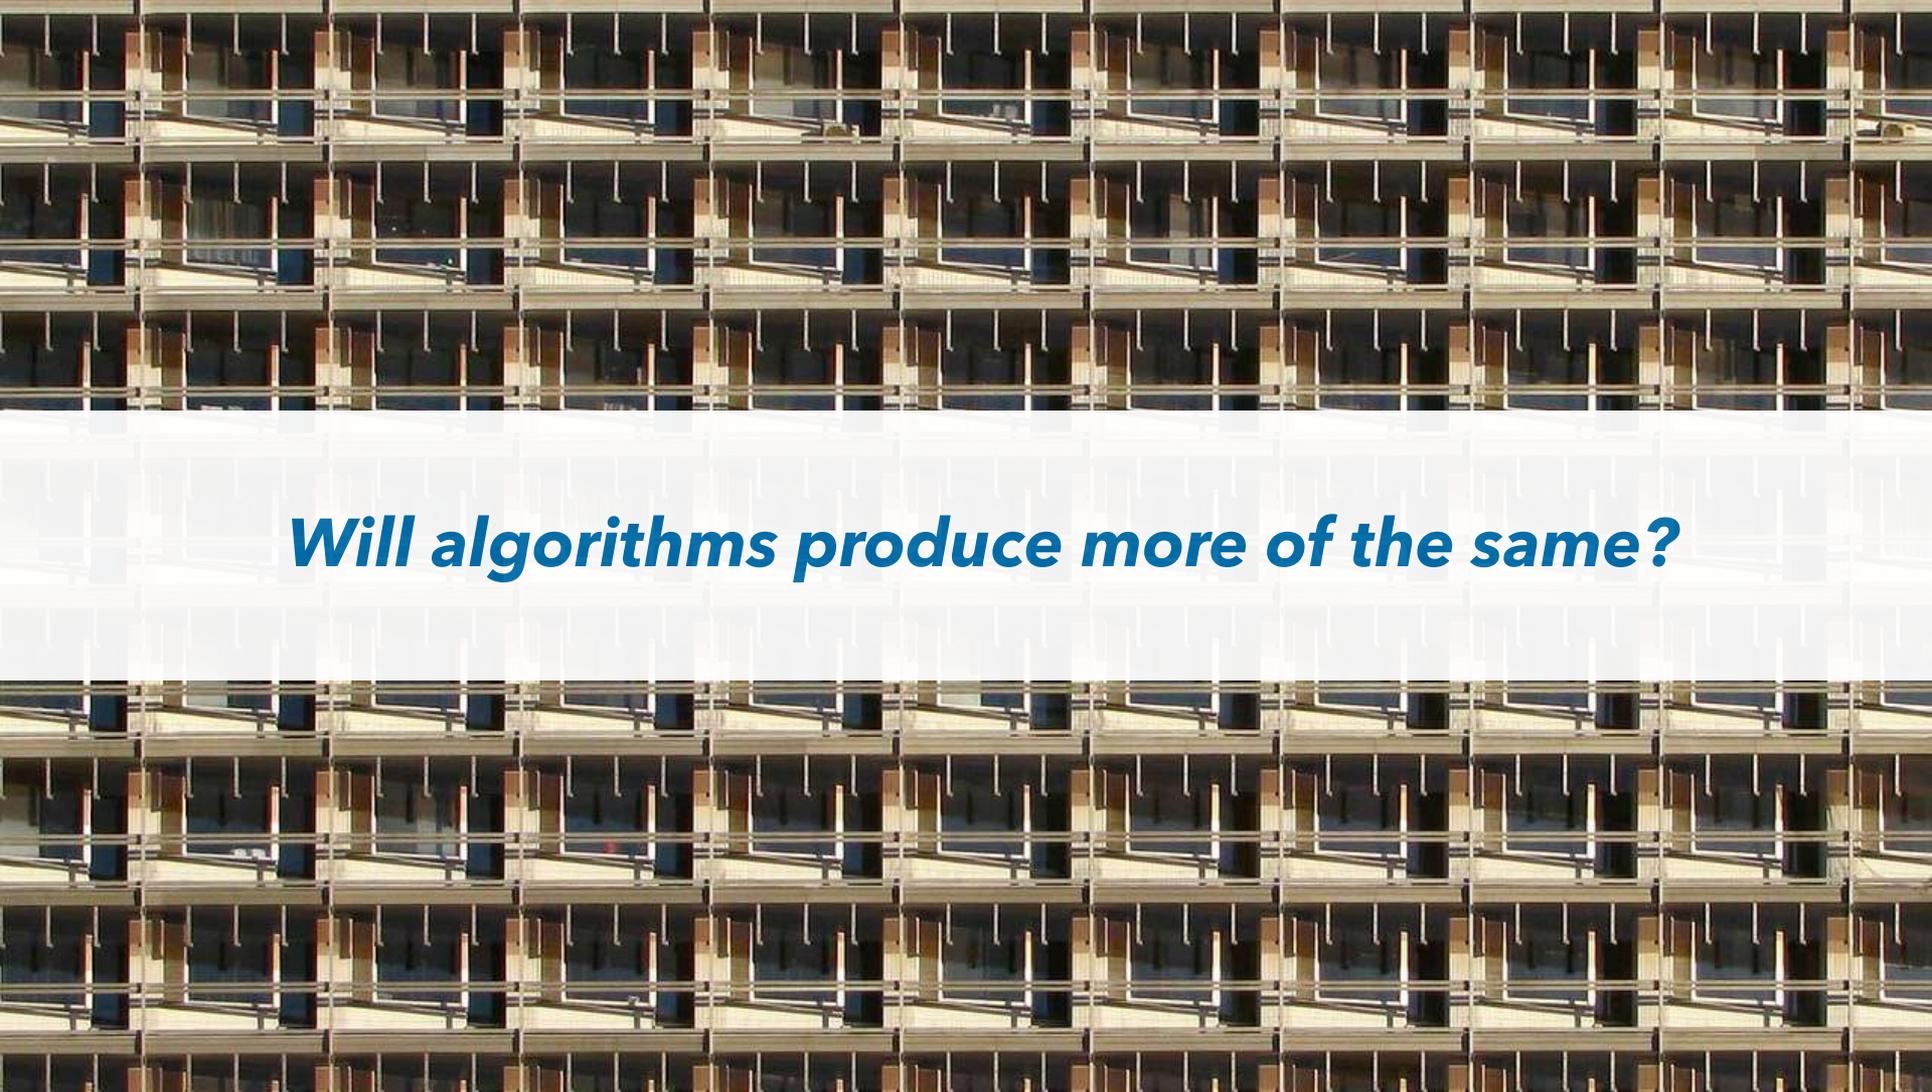
Abspielposition
 Mikrofonpegel (Aufnahmen in rot)
 Aufnahmeposition

Reflection



There are many challenging questions left open.





Will algorithms produce more of the same?

The background of the slide features a warm, orange-toned sunset. In the center, a person is silhouetted with their arms raised in a gesture of triumph or joy. Below them, a dog is also silhouetted, looking up towards the person. In the bottom left corner, a bright sun is partially visible. In the bottom right corner, two birds are silhouetted in flight against the sky. A white horizontal band runs across the middle of the image, containing the text.

Will music creators be satisfied with what they deliver?

A silhouette of a person with long hair, standing on a wooden pier or boardwalk. Their arms are outstretched horizontally, and their head is tilted back, looking towards the sky. The background is a bright, golden sunset over the ocean, with the sun low on the horizon. The sky transitions from a pale yellow near the horizon to a soft purple and pink at the top. The water in the foreground is dark with some white foam from waves.

Will music consumers be satisfied with what they get?



Will there be a rush of self-appointed creators?



The image features a dense field of dark, irregularly shaped pebbles or stones. The colors range from deep black to dark grey, with some lighter, brownish-grey spots scattered throughout. In the lower-middle section, a single, bright blue pebble stands out prominently, drawing the viewer's eye. The overall texture is rough and granular.

Will it be easy to identify the good creations?



***Will music lose its value?
Will human artists lose their value?***





***What about copyright / right d'auteur?
Who is the creator / author?***



Various stakeholders are affected!

end consumers

music creators

performers

music companies

platform providers

the society

European
Platform for
Digital
Humanism

Out of the Box

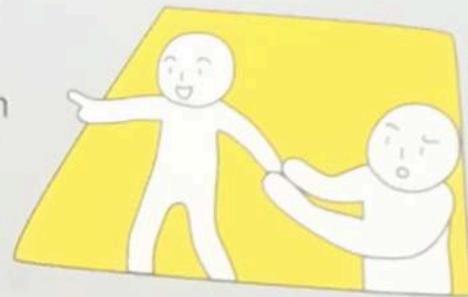
The Midlife Crisis of the Digital Revolution

ARS ELECTRONICA

Festival for Art, Technology and Society

POSTCITY Linz

September 5–9, 2019



AI x Music Festival

Encounters in the uncharted territories between human creativity and mechanical perfection

Das Festival AIxMusic (AI trifft Musik) wird von Ars Electronica und der Europäischen Kommission im Rahmen der Initiative STARTS in Kooperation mit vielen Partnern aus Industrie, akademischer Forschung sowie Kunst- und Bildungseinrichtungen organisiert. Es ist ein Hybrid aus einer Präsentation von Start-ups und einer Präsentation von Start-ups. Es beschäftigt sich mit KI und Musik, um die Herausforderungen der KI-Forschung im Hinblick auf ihre möglichen Auswirkungen auf unser Leben zu demonstrieren und zu diskutieren.

Regarding Digital Humanism...

Musik trifft Technologie trifft Musik

Kreativität – die Fähigkeit, neue und unerwartete Ausdrucksformen hervorzubringen, über die Wiederholung und Veränderung des Bestehenden hinauszugehen – wird oft als ultimative Grenze angesehen, die nur dem menschlichen Geist möglich ist.

- Werden Maschinen in der Lage sein, überzeugende Kunstwerke oder wissenschaftliche Theorien zu schaffen?



Vienna Manifesto on Digital Humanism

A call for action!

Digital Humanism refers to an approach that describes, analyzes, and influences the complex interplay of technology and humankind. Its goal is to **put the human at the center of technological progress** to shape current and future technologies.



Vienna Manifesto on Digital Humanism

...promote **democracy and inclusion** ... safeguard **free expression** ... regulations should **ensure fairness, accountability, and transparency** of algorithms ... **intervene with tech monopolies** ... **breaking disciplinary silos** ... **engage** with the wider society ... **take responsibility** ... technology is not neutral ... a vision is needed for **new educational curricula** ...

...we must go into action and take the right direction!

VIENNA MANIFESTO ON DIGITAL HUMANISM

VIENNA, MAY 2019

»**The system is failing**« – stated by the founder of the Web, Tim Berners-Lee – emphasizes that while digitalization opens unprecedented opportunities, it also raises serious concerns: the monopolization of the Web, the rise of extremist opinions and behavior orchestrated by social media, the formation of filter bubbles and echo chambers as islands of disjoint truths, the loss of privacy, and the spread of digital surveillance. Digital technologies are disrupting societies and questioning our understanding of what it means to be human. The stakes are high and the challenge of building a just and democratic society with humans at the center of technological progress needs to be addressed with determination as well as scientific ingenuity. Technological innovation demands social innovation, and social innovation requires broad societal engagement.

This manifesto is a call to deliberate and to act on current and future technological development. We encourage our academic communities, as well as industrial leaders, politicians, policy makers, and professional societies all around the globe, to actively participate in policy formation. Our demands are the result of an emerging process that unites scientists and practitioners across fields and topics, brought together by concerns and hopes for the future. We are aware of our joint responsibility for the current situation and the future – both as professionals and citizens.

Today, we experience the co-evolution of technology and humankind. The flood of data, algorithms, and computational power is disrupting the very fabric of society by changing human interactions, societal institutions, economies, and political structures. Science and the humanities are not exempt. This disruption simultaneously creates and threatens jobs, produces and destroys wealth, and improves and damages our ecology. It shifts power structures, thereby blurring the human and the machine.

The quest is for enlightenment and humanism. The capability to automate human cognitive activities is a revolutionary aspect of computer science / informatics. For many tasks, machines surpass already what humans can accomplish in speed, precision, and even analytic deduction. The time is right to bring together humanistic ideals with critical thoughts about technological progress. We therefore link this manifesto to the intellectual tradition of humanism and similar movements striving for an enlightened humanity.

Like all technologies, digital technologies do not emerge from nowhere. They are shaped by implicit and explicit choices and thus incorporate a set of values, norms, economic interests, and assumptions about how the world around us is or should be. Many of these choices remain hidden in software programs implementing algorithms that remain invisible. In line with the renowned Vienna Circle and its contributions to modern thinking, we want to espouse critical rational reasoning and the interdisciplinarity needed to shape the future.

We must shape technologies in accordance with human values and needs, instead of allowing technologies to shape humans. Our task is not only to rein in the downsides of information and communication technologies, but to encourage human-centered innovation. We call for a Digital Humanism that describes, analyzes, and, most importantly, influences the complex interplay of technology and humankind, for a better society and life, fully respecting universal human rights.

In conclusion, we proclaim the following core principles:

- **Digital technologies should be designed to promote democracy and inclusion.** This will require special efforts to overcome current inequalities and to use the emancipatory potential of digital technologies to make our societies more inclusive.
- **Privacy and freedom of speech are essential values for democracy and should be at the center of our activities.** Therefore, artifacts such as social media or online platforms need to be altered to better safeguard the free expression of opinion, the dissemination of information, and the protection of privacy.
- **Effective regulations, rules and laws, based on a broad public discourse, must be established.** They should ensure prediction accuracy, fairness and equality, accountability, and transparency of software programs and algorithms.
- **Regulators need to intervene with tech monopolies.** It is necessary to restore market competitiveness as tech monopolies concentrate market power and stifle innovation. Governments should not leave all decisions to markets.
- **Decisions with consequences that have the potential to affect individual or collective human rights must continue to be made by humans.** Decision makers must be responsible and accountable for their decisions. Automated decision making systems should only support human decision making, not replace it.
- **Scientific approaches crossing different disciplines are a prerequisite for tackling the challenges ahead.** Technological disciplines such as computer science / informatics must collaborate with social sciences, humanities, and other sciences, breaking disciplinary silos.
- **Universities are the place where new knowledge is produced and critical thought is cultivated.** Hence, they have a special responsibility and have to be aware of that.
- **Academic and industrial researchers must engage openly with wider society and reflect upon their approaches.** This needs to be embedded in the practice of producing new knowledge and technologies, while at the same time defending the freedom of thought and science.
- **Practitioners everywhere ought to acknowledge their shared responsibility for the impact of information technologies.** They need to understand that no technology is neutral and be sensitized to see both potential benefits and possible downsides.
- **A vision is needed for new educational curricula, combining knowledge from the humanities, the social sciences, and engineering studies.** In the age of automated decision making and AI, creativity and attention to human aspects are crucial to the education of future engineers and technologists.
- **Education on computer science / informatics and its societal impact must start as early as possible.** Students should learn to combine information-technology skills with awareness of the ethical and societal issues at stake.

We are at a crossroads to the future; we must go into action and take the right direction!

VIENNA MANIFESTO ON DIGITAL HUMANISM

VIENNA, MAY 2019

The system is falling apart – created by the founder of the WWW, Tim Berners-Lee – emphasizes that while digitalization opens unprecedented opportunities, it also raises serious concerns: the monopolization of the data for the extreme opinions and behaviors, exclusion by the algorithms in order of the hidden content of the data, islands of digital reefs, inclusion of the data and the spread of digital surveillance. Digital technologies are changing how we live and work, but our understanding of what it means to be human. The stakes are high: building a just and democratic society with humans at the center. The challenges that need to be addressed with determination as well as scientific rigor and care. The future demands social innovation and social innovation requires



Support the Manifesto as an individual or organization and help us spread the word

<https://www.informatik.tuwien.ac.at/dighum/>

Like all technologies, digital technologies do not emerge from nowhere. They are shaped by implicit and explicit choices and thus incorporate a set of values, norms, economic interests, and assumptions about how the world around us is or should be. Many of these choices remain hidden in software programs implementing algorithms that remain invisible. In line with the renowned Vienna Circle and its contributions to modern thinking, we want to espouse critical rational reasoning and the interdisciplinarity needed to shape the future.

We must shape technologies in accordance with human values and needs, instead of allowing technologies to shape humans. Our task is not only to rein in the downsides of information and communication technologies, but to encourage human-centered innovation. We call for a Digital Humanism that describes, analyzes, and, most importantly, influences the complex interplay of technology and humankind, for a better society and life, fully respecting universal human rights.

In conclusion, we proclaim the following core principles:

- **Digital technologies should be designed to promote democracy and inclusion.** This will require special efforts to overcome current inequalities and to use the emancipatory potential of digital technologies to make our societies more inclusive.
- **Privacy and freedom of speech are essential values for democracy and should be at the center of our activities.** Therefore, artifacts such as social media or online platforms need to be altered to better safeguard the free expression of opinion, the dissemination of

- **Effective regulations, rules and laws, based on a broad public discourse, must be established.** They should ensure protection of privacy, fairness and equality, accountability, and transparency of software programs and algorithms.

- **Regulators need to intervene with tech monopolies.** It is necessary to restore market competitiveness as tech monopolies concentrate market power and stifle innovation. Governments should not leave all decisions to markets.

- **Universities are the place where new knowledge is produced and critical thought is cultivated,** hence they have a special responsibility and have to be aware of that.

- **Academic and industrial researchers must engage openly with wider society.**

- **Practitioners everywhere ought to acknowledge their shared responsibility for the impact of information technologies.** They need to understand that no technology is neutral and be sensitized to see both potential benefits and possible downsides.

- **A vision is needed for new education programs, combining knowledge from the humanities, the social sciences, and engineering studies.** In the age of automated decision making and AI, creativity and attention to human aspects are crucial to the education of future engineers and technologists.

- **Education on computer science / informatics and its societal impact must start as early as possible.** Students should learn to combine information-technology skills with awareness of the ethical and societal issues at stake.

We are at a crossroads to the future; we must go into action and take the right direction!

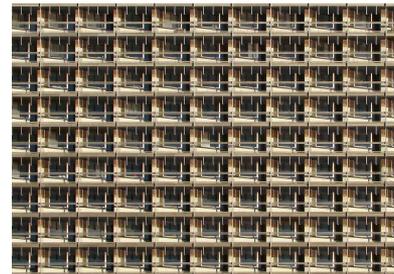
Summing up



New ideas



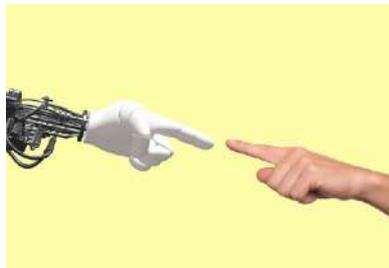
Getting out of the bubble



More of the same



Rush of creators



Human-computer symbiosis



Filtering noise

Wrap up



Many open questions

Speakers

Christine Bauer

Johannes Kepler University Linz

✉ christine.bauer@jku.at
🐦 [@christine_bauer](https://twitter.com/christine_bauer)



Peter Knees

TU Wien

✉ peter.knees@tuwien.ac.at
🐦 [@peter_knees](https://twitter.com/peter_knees)



Richard Vogl

TU Wien

✉ richard.vogl@tuwien.ac.at
🐦 [@richard_vogl](https://twitter.com/richard_vogl)



Hansi Raber

✉ super@superduper.org
🐦 [@kritzikratzi](https://twitter.com/kritzikratzi)
📺 <https://www.youtube.com/hansiraber>

