Scientia
AUGUST 2025

Professor Jinhee Choi
Dr Jaeseong Jeong
Mr DongHyeon Kim

Revolutionising Chemical Safety: How AI Could Replace Animal Testing

doi.org/10.33548/SCIENTIA1231

ENGINEERING & COMPUTER SCIENCE

Modern life exposes us to a staggering array of synthetic substances—over 350,000 chemicals are registered for use worldwide, found in products as diverse as pesticides, plastics, cosmetics, and pharmaceuticals. Ensuring the safety of all these compounds is a daunting task. Researchers led by Prof Jinhee Choi at University of Seoul are developing cutting-edge artificial intelligence models to predict the potential toxicity of chemicals, with the aim of reducing the need for controversial animal experiments. By analysing vast toxicology databases and pioneering new AI techniques, Prof Choi's team is working towards a future where the safety of everything from industrial compounds to household products can be assessed more quickly, cheaply, and humanely.

Cite as SD/Choi/Revolutionising Chemical Safety: How AI Could Replace Animal Testing/August 2025/1231

The Problem with Animal Testing

Traditionally, the toxicity of chemicals has been evaluated through animal testing. In these experiments, creatures like mice, rats, rabbits, and even dogs are exposed to substances, to see if they cause adverse effects such as illness, reproductive problems, or abnormalities in offspring. It's an approach that dates back to the early 20th century and has long been considered the gold standard in toxicology. However, as Prof Jinhee Choi explains, animal testing has significant drawbacks.

The limitations of animal testing are particularly evident in the field of chemical safety assessment. A single traditional toxicity study can require hundreds of animals and take several years to complete, costing hundreds of thousands of pounds. This creates a significant bottleneck in chemical safety assessment, as thousands of new chemicals enter commerce each year. These challenges have become increasingly pressing as regulations worldwide demand more comprehensive safety data for chemicals. For example, the Korea and European Union's REACH regulations require safety information for all chemicals produced or imported in quantities over one tonne per year. Meeting these requirements through traditional animal testing would be practically impossible, requiring millions of animals and billions in funding.

Perhaps most concerning from a scientific standpoint, results from animal tests do not always translate well to humans. There are numerous cases of substances that were deemed safe based on animal studies but later turned out to cause problems in people, such as thalidomide. Conversely, some promising drugs have been abandoned because they caused side effects in animals that may not be relevant to humans. More fundamentally, there are growing ethical concerns about subjecting sentient creatures to potentially toxic substances for the sake of human safety. Many countries have passed laws to reduce animal testing, and public opinion is increasingly opposed to the practice.

Prof Choi and her team are working to address these challenges by developing new ways to assess chemical safety without relying on animal testing. The team's core innovation lies in their strategy of using toxicity big data qand AI, within an adverse outcome pathway (AOP) framework that makes the results accurate and explainable. On one side, they work with models based on in vitro data – laboratory tests performed on cells or biological molecules. These come primarily from the ToxCast database, which contains millions of data points showing how chemicals interact with various biological processes. On the other side, they have models based on in vivo toxicity data – results from historical animal studies that show how chemicals cause harm in living organisms.

Big Data in Toxicology

One of the key resources in this effort is the ToxCast database, a massive collection of chemical testing data maintained by the US Environmental Protection Agency. Prof Choi's team has conducted an extensive analysis of this database, which contains results from over 1,500 different types of tests performed on more than 10,000 chemicals. These tests examine how chemicals interact with various biological processes in cells. For example, some tests examine whether chemicals interfere with important cellular proteins called nuclear receptors, while others examine the effects on enzymes called kinases that help control cell behaviour. This information helps scientists understand how chemicals might cause toxicity without needing to test them in animals.

The team carried out a comprehensive review of how researchers are using this wealth of data. Their analysis of 96 research papers revealed that scientists are increasingly turning to machine learning

approaches to predict chemical toxicity, with such studies increasing five-fold between 2014 and 2020. This dramatic increase reflects both the growing availability of high-quality data and improvements in computational methods. The team's analysis revealed that the ToxCast database is particularly valuable because it provides standardised, high-quality data across many different types of biological processes. This standardisation is crucial for developing reliable predictive models. The database includes information about how chemicals interact with various cellular components, from proteins that control gene expression to enzymes that break down toxic substances.

Understanding this data requires careful consideration of how cells respond to chemicals. For example, some chemicals might appear toxic simply because they kill cells at high concentrations rather than through specific biological mechanisms. The team's analysis helped identify ways to distinguish between these general toxic effects and more specific interactions that might indicate particular types of toxicity.

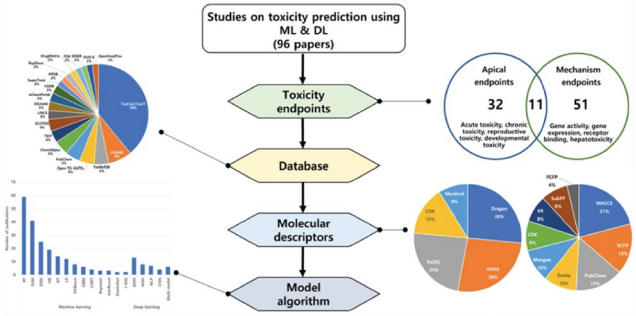

The Power of Artificial Intelligence

Building on this foundation, the team developed sophisticated machine learning models to predict how chemicals might affect different biological processes. These models essentially teach computers to recognise patterns in chemical structures that might make them toxic. The team's approach to machine learning is particularly innovative in its combination of different data types. Rather than simply looking at chemical structures, their models incorporate information about how chemicals interact with biological systems. This helps the models make more accurate predictions and provides insights into the mechanisms by which chemicals might cause toxicity.

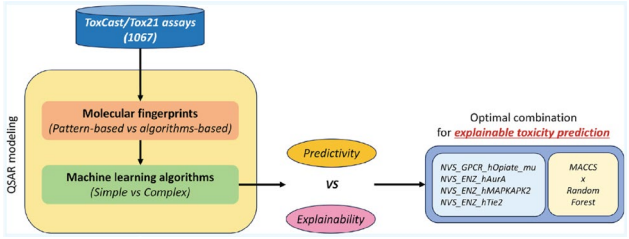

The process starts by representing each chemical in a way that a computer can understand. One common approach is to create a 'molecular fingerprint' – a digital representation that encodes key features of a chemical's 2D structure, such as the presence of certain atoms, functional groups, or ring systems. Prof Choi likens it to a barcode that uniquely identifies each compound.

With the chemicals digitised in this way, the next step is to train machine learning algorithms to find patterns in the data. The basic idea is to show the algorithm examples of chemicals with known toxicity, represented by their fingerprints, and let it figure out what structural features are associated with harmful effects. This is a bit like teaching a child to recognise animals. You might show them pictures of different creatures, pointing out features like fur, feathers, beaks, or paws, and labelling each one as a mammal, bird, reptile, etc. With enough examples, the child starts to learn the patterns – furry things with four legs are usually mammals, feathered creatures with wings are generally birds, and so on. They can then use this knowledge to make educated guesses about new animals they have not seen before. Similarly, Prof Choi's models learn to associate certain chemical features with toxicity. Once trained, they can be used to predict the likely effects of new, untested chemicals based on their structural similarity to known toxicants.

The researchers found that simpler machine learning algorithms, particularly a method called Random Forest, often work better than more complex approaches when it comes to balancing accuracy with explainability (the ability to understand why the model made particular predictions). They also carefully chose how to represent chemical structures in their models, using a molecular fingerprint called MACCS keys that captures meaningful chemical features that scientists can interpret. This focus on explainability is crucial because safety assessments need to be trustworthy and transparent. While some artificial intelligence approaches act as 'black boxes', making predictions without clear explanations, Prof Choi's team specifically developed models that can explain their reasoning. This makes them more useful for regulatory purposes, where decisions about chemical safety need to be based on clear evidence.

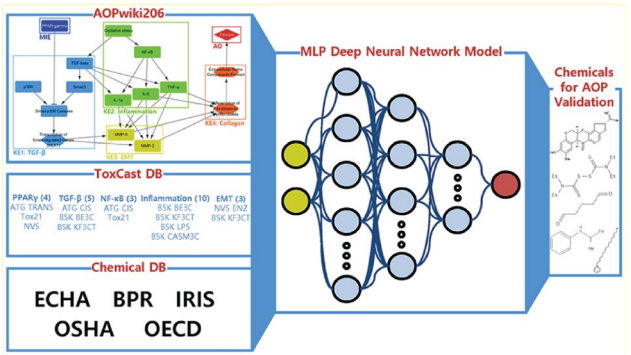

Integration with the Adverse Outcome Pathway (AOP) Framework

One of the most innovative aspects of the team's work is their integration of artificial intelligence with what scientists call the Adverse Outcome Pathway (AOP) framework. An AOP is essentially a flowchart that describes, step by step, how exposure to a chemical leads to an adverse health effect. The AOP framework breaks down complex biological processes into a series of connected events. It starts with what scientists call a molecular initiating event – the first interaction between a chemical and a biological molecule. This triggers a series of key events, each representing a change in biological function, ultimately leading to adverse outcomes – the actual harm to health.

In a case study demonstrating this approach, the team focused on how certain chemicals might cause pulmonary fibrosis – a serious and often irreversible lung disease characterised by scarring of the lung tissue. Pulmonary fibrosis can be caused by inhaling certain chemicals, particularly in industrial settings. The disease begins when these substances bind to and inhibit a protein in lung cells called peroxisome proliferator-activated receptor gamma (PPARγ). This triggers a complex chain of events, including inflammation, changes in cell behaviour, and excessive deposition of extracellular matrix proteins like collagen, gradually leading to scar tissue buildup and lung function loss.

The team identified relevant biological assays in the ToxCast database corresponding to key events in the progression from PPARγ

inhibition to fibrosis to model this pathway. For example, they found assays measuring the activity of inflammatory signalling proteins, the expression of genes involved in the extracellular matrix, and markers of epithelial-mesenchymal transition. In this process, lung cells start to take on characteristics of scar-forming fibroblasts. For chemicals not covered in the ToxCast database, the team used their previously developed machine learning models to predict whether the substances would be active in these assays based on their structural similarity to known PPARγ inhibitors and other fibrosis-inducing compounds.

By combining their artificial intelligence models with this pathway-based understanding, the team created a multi-layer model that could predict not only a chemical's likelihood of causing pulmonary fibrosis but also its potential to trigger each key step in the AOP. The potential benefits of this integrated approach are significant. By combining in vitro data models, which are good at identifying initial biological interactions, with in vivo data models that capture real-world toxic effects, the team hopes to create a system that can make more accurate and scientifically defensible predictions. The AOP framework strengthens model credibility by aligning computational predictions with well-established biological pathways, enhancing their utility for regulatory decision-making..

Practical Applications in Modernising Chemical Risk Assessment

While impressive from a technical standpoint, the ultimate test of any toxicity prediction model is whether it can be used to assess the safety of chemicals in the real world. This is especially important for substances already in widespread use that have not undergone extensive safety testing, such as those found in many consumer products. To address this challenge, Prof Choi's team have developed models specifically for predicting inhalation toxicity of everyday items like air fresheners, cleaning products, and personal care products. These types of household goods are ubiquitous, but surprisingly little is known about the inhalation toxicity of many of their ingredients.

The team built these models using data from standardised safety tests collected from an international database. When applied to 689 chemicals used in consumer products, their models identified 79 chemicals of potential concern. Notably, many of these chemicals currently lack proper safety classifications, highlighting potential gaps in current regulatory systems. This work demonstrates how artificial intelligence can help address real-world chemical safety challenges.

The models can quickly screen large numbers of chemicals, identifying those that might need more detailed safety assessment. This is particularly valuable given the vast number of chemicals in use and the limited resources available for safety testing.

The team's approach is especially innovative in its use of multiple types of data and models. They combined information from traditional animal studies with results from cell-based tests and computer predictions. This integrated approach helps provide a more complete picture of chemical safety than any single method alone. The researchers also developed ways to determine when their models' predictions might be less reliable. This is crucial for practical applications, as it helps users understand when additional testing might be needed. The team's work showed that their models were particularly good at identifying potentially harmful chemicals, though they sometimes flagged safe chemicals as potentially dangerous.

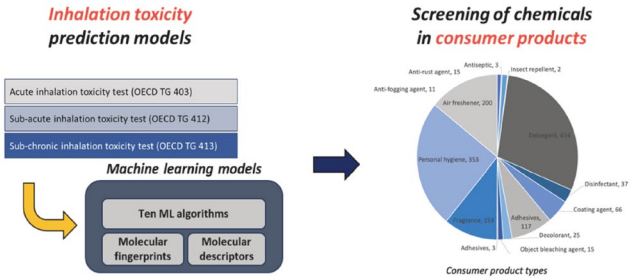

Challenges and Future Directions

While these advances are promising, challenges remain. The team acknowledges that their current models sometimes err on the side of caution, flagging chemicals as potentially harmful when they might be safe. While this is preferable to missing truly harmful chemicals, it means additional testing is still needed to confirm the models' predictions.

Another major challenge is the limited availability of high-quality data for many types of toxic effects. While databases like ToxCast provide extensive information about how chemicals interact with cellular processes, there is often less data available about how these interactions lead to actual harm in living organisms. This makes it difficult to fully validate predictive models. A further challenge lies in understanding the complex relationships between different biological processes.

While the AOP framework helps map out how chemicals might cause harm, many toxic effects involve multiple interacting pathways. Developing models that can account for these complex interactions remains a significant challenge.

Looking ahead, Prof Choi and her colleagues continue to refine their approach, working to improve both the accuracy of their predictions and the clarity of their explanations. Their work demonstrates how careful integration of different types of data and models, guided by biological understanding, can help solve one of the major challenges in artificial intelligence – creating systems that not only make accurate predictions but can also explain their reasoning in scientifically meaningful ways.

The team has already presented their approach at several international conferences and submitted it as a new approach methodology (NAM) to the European Partnership for Alternative Approaches to Animal Testing. They are currently working to publish their integrated framework, which could represent a significant advance in the field of computational toxicology. If successful, their approach could help reduce reliance on animal testing while providing more transparent and scientifically rigorous ways to assess chemical safety.

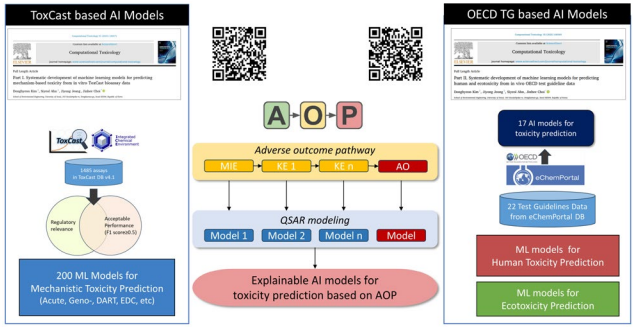

MEET THE RESEARCHERS

Professor Jinhee Choi
School of Environmental Engineering,
University of Seoul, Korea

Dr Jinhee Choi is a Professor at the School of Environmental Engineering, University of Seoul, and Director of the Chemical Big Data AI Research Center. She is an internationally recognized expert in computational toxicology, AI-driven toxicity prediction, systems toxicology, and Next-Generation Risk Assessment (NGRA). Her research focuses on developing explainable AI models and AOP-based strategies to modernize chemical safety assessment, including predictive frameworks for chemical mixtures and regulatory applications.

Prof Choi's work highlights bridging human toxicology and ecotoxicology under a One Health approach, aligning AI-powered models with AOPs to improve hazard identification and chemical prioritization. Recognized as a Highly Cited Researcher in Toxicology and one of the World's Top 2% Scientists, she has also served as President of the Korean Society of Environmental Toxicology and Health and as a member of the Korean Presidential Advisory Council on Science and Technology.

Through her pioneering research and leadership, Prof Choi continues to advance the development of safer chemicals and sustainable risk assessment strategies in the era of next-generation regulatory science.

CONTACT

jinhchoi@uos.ac.kr

https://est.uos.ac.kr/

linkedin.com/in/jinhee-choi-0aa2b43b

Chemical Bigdata AI Center chemBAI

ChemBAi Chemical Bigdata AI Solution

ESTOX Environmental Systems Toxicology Lab

Dr Jaeseong Jeong
School of Environmental Engineering,
University of Seoul, Korea

Dr Jaeseong Jeong is based in the School of Environmental Engineering at the University of Seoul. His research focuses on developing adverse outcome pathway networks, toxicity prediction models combining toxicity databases and artificial intelligence, and new approach methodologies for risk assessment. Dr Jeong received his PhD from the University of Seoul, where his thesis focused on developing an adverse outcome pathway network for prioritising chemicals using computational approaches. He has published 31 peer-reviewed papers that have been cited over 1300 times. In 2022, Dr Jeong received the Young Scientist Award from the Korean Society of Environmental Health and Toxicology.

CONTACT

erphios@naver.com

https://www.linkedin.com/in/jseongjeong

Mr DongHyeon Kim
School of Environmental Engineering,
University of Seoul, Korea

Donghyeon Kim is a PhD student in Environmental Engineering at the University of Seoul, Korea, and in Bioinformatics at Université Paris Cité, France. His research interests include developing AI-based toxicity prediction models, omics, adverse outcome pathways (AOPs), and bridging toxicology with epidemiology. Mr. Kim has published 15 peer-reviewed papers during his PhD program. He has been awarded the Presidential Science Scholarship by the Ministry of Science and ICT (MSIT) and the 3R Student Grant Award by the EU JRC·EPAA.

CONTACT

kdhyeon0096@gmail.com

linkedin.com/in/donghyeon-kim-81a5a32bb

FURTHER READING

D. Kim, J. Choi, *AI-driven hazard prioritization of plastic additives using Tox21*, Environmental Toxicology & Chemistry, 2025, DOI: https://doi.org/10.1093/etojnl/vgaf228

D. Kim, S. Ahn, J. Choi, *Identification of developmental and reproductive toxicity of biocides in consumer products using ToxCast bioassays data and machine learning models*, Environment International, 2025, DOI: https://doi.org/10.1016/j.envint.2025.109621

D. Kim, J. Choi, *AI-based Toxicity Prediction Models using ToxCast Data: Current Status and Future Directions for Explainable Models*, Toxicology, 2025, DOI: https://doi.org/10.1016/j.tox.2025.154230 ResearchGate

D. Kim, S. Cho, J.-J. Jeon, J. Choi, *Inhalation Toxicity Screening of Consumer Products Chemicals using OECD Test Guideline Data-based Machine Learning Models*, Journal of Hazardous Materials, 2024, DOI: https://doi.org/10.1016/j.jhazmat.2024.135446

J Jeong, D Kim, J Choi, Application of ToxCast/Tox21 data for toxicity mechanism-based evaluation and prioritization of environmental chemicals: Perspective and limitations, *Toxicology in Vitro*, 2022, DOI: https://doi.org/10.1016/j.tiv.2022.105451

J Jeong, J Choi, Artificial Intelligence-Based Toxicity Prediction of Environmental Chemicals: Future Directions for Chemical Management Applications, *Environmental Science & Technology*, 2022, DOI: https://doi.org/10.1021/acs.est.1c07413

J Jeong, N Garcia-Reyero, L Burgoon, *et al.*, Development of Adverse Outcome Pathway for PPARγ Antagonism Leading to Pulmonary Fibrosis and Chemical Selection for Its Validation: ToxCast Database and a Deep Learning Artificial Neural Network Model-Based Approach, *Chemical Research in Toxicology*, DOI: https://doi.org/10.1021/acs.chemrestox.9b00040

D. Kim, J. Jeong, S. Ahn, J. Choi, Part I. Systematic development of machine learning models for predicting mechanism-based toxicity from in vitro ToxCast bioassay data, *Computational Toxicology*, 2025, DOI: https://doi.org/10.1016/j.comtox.2025.100371

D. Kim, J. Jeong, S. Ahn, J. Choi, Part II. Systematic development of machine learning models using toxicity data generated in accordance with OECD Test Guidelines (TG), *Computational Toxicology*, 2025, DOI: https://doi.org/10.1016/j.comtox.2025.100369