

THE RESEARCH UNIVERSITY PRESS

International
Computational Law Journal

VOLUME 2 / ISSUE 5 / CURATED SOURCE EDITION

Rapid Prototyping of Legal Tech and Design Using Generative AI: Bridging Theory and Practice in Legal Education

Barbora Obračajová

MIT Computational Law Report · Originally published 2024-09-10

Independent reader edition under CC BY 4.0. This attribution cover was added; original source PDF pages follow unchanged. Original authorship and source identity are retained. MIT, Stanford University, Stanford Law School and CodeX have not sponsored or endorsed this republication.

Source, license and version evidence accompany the online article.

MIT Computational Law Report

Rapid Prototyping of Legal Tech and Design Using Generative AI: Bridging Theory and Practice in Legal Education

Barbora Obračajová

Published on: Sep 10, 2024

URL: https://law.mit.edu/pub/rapidprototypingoflegaltech

License: Creative Commons Attribution 4.0 International License (CC-BY 4.0)

ABSTRACT

We are collectively navigating the discovery stage of working with Generative Artificial Intelligence ("Generative AI") and Large Language Model Applications ("LLM Applications"). This phase comes with unique challenges and opportunities for legal education and practice. This article explores a series of experiments conducted in the Modern Lawyers Course at the Faculty of Law, Charles University in Prague, which utilized the Rapid Prototyping Method with Generative AI to create high-fidelity, low-commitment prototypes that allow for quick validation of legal tech designs. This approach has not only enhanced the discovery process and legal tech development but also equipped a new generation of lawyers with key technology skills and new legal problem-solving abilities.

Introduction

When was the last time you started tinkering with something new in your legal work? Sketched out a new legal tech bot or drew out the proceedings as a map? Chances are that lately, it has been a little bit more frequent thanks to the general excitement over the capabilities of the Generative Artificial Intelligence ("Generative AI"). We are at the discovery stage, collectively uncovering the possibilities, limitations, and most fitting use cases for the Large Language Models ("LLM") and their applications ("LLM Applications").

There is no settled range of best practices - yet. The future is up to us to shape, and this will require active ventures into uncharted territory. It won't be perfect right from the start. To truly understand what we are dealing with, we need to get our hands dirty and encounter many dead ends. In this article, we will discuss one of such quests toward practical, hands-on legal prototyping, supercharged by Generative AI.

The following is a reflection of a series of experiments performed within the legal design course Modern Lawyers taught by the author at the Faculty of Law, Charles University in Prague, Czech Republic.

The key learning from these experiments is that Generative AI can provide easier access to relatively high-fidelity yet low-commitment prototypes that can be used for the validation of ideas for legal design or tech solutions. Generative AI has the potential not only to enhance legal tech development but also to empower a new generation of lawyers equipped with a blend of legal and tech-savvy skills and to introduce a new vision for approaching tech stack ideation and procurement.

This article first describes the general context and methodology of rapid prototyping within the Modern Lawyers class, then delves into three specific examples of students using Generative AI to prototype in novel ways. To generalize the argument, this article will focus on visualizations, chatbots, and service maps, which are central to the curriculum.

The observations are not conclusive but represent potential explorations within the field. Issues related to copyright, data protection, and liability for AI-generated advice are beyond the scope of this article.

Using Generative AI for Rapid Prototyping

Prototyping can be considered controversial among lawyers. The broad narrative of the industry is often geared towards early hyper-specialization and little tolerance for experimentation. Neither lends itself well to noncommittal exploration of new concepts. On top of that, the legal rules are complex and the service delivery models are predominantly rigid, which complicates the possibility of translating law into any kind of product. This can limit breakthrough ideas by drowning them in the text to be parsed, excessive focus on finding exceptions, or in a rather difficult change management process. All these factors quite often push design considerations to the sidelines.

Generative AI provides an alternative approach to these roadblocks in the form of rapid prototyping. In short, thanks to the text-processing possibilities of LLM Applications, we can now conceptualize a potential mockup for any tool or rendition of legal information much quicker. This makes prototyping much swifter and readily available while also moving the focus of the legal founders from going through loads of legal text to crafting a brilliant user experience.

Rapid prototyping is an integral part of the Modern Lawyers course curriculum, which we will discuss in the next sections.

Modern Lawyers Course

The Modern Lawyers course has been taught at the Faculty of Law, Charles University in Prague since 2022. Its goal is to empower the alumni to explore alternative approaches to solving legal problems and apply creative solutions to meet legal needs.

Over the course of the semester, the students explore the theoretical background of legal design, legal operations, legal project management, process design, inclusive design, ethics, and systems design, alongside plenty of hands-on prototyping experience. Learning is facilitated within a series of assignments relating to the context discussed within the classes, each more challenging than the last.



Figure 1: Modern Lawyers Course Map Outlining the Key Topics for Each Class

Based on these experiments, the students ideate and develop their final project, including the main logic, processes, and prototypes. Each student team works with an expert mentor from the field - ranging from seasoned legal professionals to developers and startup founders.

In each final project, the students explore a multitude of ways of building with and without the help of Generative AI that would each require a lot more space to describe and appreciate. Thus this article focuses only on the core curriculum exercises.

The course design is iterative and data-driven, so there have been different approaches applied in each semester, also in light of the developing technological advancements.

Rapid Prototyping in the Course

Instead of being treated as separate topics, Generative AI and prompt engineering are integrated and applied throughout the course. In the context of the challenges and tasks for each consecutive class, students work with text, actively using Generative AI to furbish and format information for their projects and cope with strict and intentionally challenging time restrictions.

The students then transform the generated outputs into high-fidelity prototypes, for example by using software such as Canva or Figma, selected specifically for their easy access and less demanding learning curve. Based on their learning objectives, some students even decide to code their solutions from scratch, often using outcomes generated by LLM Applications.

Course Outcomes

Throughout the course, the students gain first-hand experience of various approaches to law and become familiar with alternative career paths in the legal industry.

Moreover, they experience both the design process as well as a critical review of their proposed solutions. The goal of these exercises is not to create something perfect or production-ready. Instead, the students test out concepts in the early stage, before the potential founders would go ahead and engage professional designers and developers.

The broader aim of these exercises is to help law students develop creative confidence and critical thinking relative to user-centricity and design decisions in law. Each student gets substantial exposure to tech, including Generative AI, as well as to practical and ethical considerations of legal design work.

The rapid prototyping also turns the process of learning the specific legal system from section and topic-based to process-based. This encourages new thinking patterns, which may mimic and reflect the potential future process of working with AI-generated legal content. Building on the inherent multidisciplinary nature of the course, this also encourages general division between front-end and back-end legal work, as well as an active approach towards testing and assumption validation, inspired by domains of design and computer science.

In the next segment, we will closely examine three selected exercises that the students undertake over the course of the semester.

Level 1: Visualizations

The Task

Upon receiving the initial orientation in the domain of legal innovation, the first practical task is to create a visualization of a document that is ubiquitous in the Czech legal system: the instructions for the parties and visitors in court hearings. This is a real document that is to be found in each courthouse and is usually the only guidance the litigants receive regarding their courtroom conduct.

The process of dissecting the legal information using Generative AI is highly idiosyncratic to the specific student teams. Below are some of the key steps that were applied in the class.

Using the Generative AI for the Dragonfly View

The first step is to take the text of the instruction and prompt an LLM Application to divide it into subsections by topic. The prompt can be along the lines of "I am going to give you a text, create an outline of its key topics." This gives the students a list of contents of the stipulations, which could form logical clusters as a basis for the visualization.

Secondly, there is the application of the language model to "translate" the text into a certain readability level, or according to the need of a specific user. Here, priming can be applied to further specify the user persona that the document is prepared for. In such a case the students would start by asking questions such as "How to ensure that a text is understandable by a 12-year-old?" and then "Rewrite the text below based on [the requirements]".

There is also space to explore more non-traditional renderings of such information instead of sticking to the traditional visualization - such as by turning the instructions into a song using music generators or by explaining the legal requirements using stories, parallels, or metaphors appropriate to the user.

The Granularity of the Legal Information

One of the early challenges is in setting the character limits for the output and making a decision regarding the level of granularity.

Lawyers (the author included) often don't stop at the first clear-cut scenario of application of the law - we are taught to explore the edge cases (see below the discussion of edge cases in chatbots). In the case of static legal information, there is always only a limited amount of how much we can fit on a page or a website. Therefore, there are inevitable design decisions as to how to capture not only the "standard" user but also alert the audience to the possibilities to diverge from the traditional path.

In such a case, rapid prototyping can very quickly explore what a certain text would look like in 200 characters, 400 characters, or more. By applying this and investigating the differences between the specific versions, the students can distill the essential but also question their and the LLM Application's heuristic judgment in which information is considered universal or crucial.

"We Are Yet to Decide as to Where We Are Going to Hang This"

One of the interesting learnings is that once the students have reached a workable prototype of their visualization, they are often confronted with the outlook for real-life practice.

Based on the inclusive design and ethics unit, the students are actively encouraged to question the accessibility and inclusiveness of their proposed solutions. Thanks to this, some of the questions students posed themselves were for example "Does it make sense to have this only as a poster hanging in a courthouse hallway when many users access information primarily online nowadays?", "Maybe we can add it to the invitation for the

court hearing so that people can prepare in advance for what is going to happen", and "If this poster is intended for children, does it make sense to hang it in the usual space - which is at the eye level of an adult?"

Using rapid prototyping, the students can explore different versions and visions for their end product, but also experience its limitations very early on in the process. This ultimately leads to more mindful design decisions, as the prototype quickly turns a broad concept into something more tangible.

Level 2: Chatbots

With law being a word-intensive discipline that often requires back-and-forth between the advisor and the client, chatbots come up quite frequently as one of the possible solutions to legal problems and design challenges. Unfortunately, the experience shows that implementing a chatbot rather often leads to impersonal and dehumanized service interactions. However, given the nature of the law, there are still opportunities for gain when using this type of interface.

The second design challenge revolves around designing a self-service solution and possibly a sales funnel for a specific legal tech solution. Students have the opportunity to develop their vision and test it. In past semesters, the applications ranged from legal intake forms for law firms, to divorce bots and apps to interact with the property registry.

The workflow is again simplified to maximize impact while optimizing efforts. The students are encouraged to take the respective laws or other legal sources and use them to prime the LLM Application. Using this pre-training, students can obtain a conversational dataset made out of the legal information that they then tweak for use in visual models using Canva, Figma Design File, or any other application depending on the tech proficiency of the respective student group.

Packaging Empathy

The first hurdle in successful chatbot development is the need to convey empathy in delivering personalized legal services. An example of this was when the students were asked to prototype a chatbot to assist the user with legal questions following the passing of a relative. This requires the chatbot to ask questions regarding the existence of will and other living loved ones in the circumstances of heightened emotional burden on the user.

Interestingly enough, illustrating these considerations required a lot of questions about the setup of the chatbot and the related crafting of the prompts. The students had a hard time deciding between (i) keeping things very matter-of-factly and serious, which resulted in unnatural, polished interaction with seemingly fake displays of sympathy, (ii) "app-ification" of emotionally charged human situations, where asking a bot regarding the passing of a loved one resembled the experience of ordering a takeout meal from a restaurant, or (iii) straightforward illustration of the purpose of the bot using somber tone while also including grim reaper to send clear expectations to the audience.

By visualizing the original idea, the students were confronted with the sheer amount of branding and design decisions that come with parsing a few paragraphs of the Civil Code into human language. Thanks to rapid prototyping, they could test various approaches while critically evaluating their prompts and the related output.

Chasing the Edge (Cases)

A chatbot is a form of dynamic legal information. That means that compared to a static poster, chatbots can be the ideal avenue to flexibly respond to users' unique information gathered from their responses. This results in the need to capture within the user journey both the most frequent uses, but also to constructively gather information from the user in case of legal exceptions and specific clauses.

With rapid GenAI prototyping, these edge cases can be explored in a multitude of ways. Firstly, there is a possibility to explicitly prompt the model to return a certain number of most frequent client situations and then explicitly ask it to generate a “non-standard” scenario that would not fit this user journey. Secondly, when the chatbot is visualized, the “non-standard” scenarios can form a part of the user testing. This can happen either directly with users' real-time participation or in a secondary way where one of the team members plays the interaction according to a prepared script, often also generated by the LLM Application. Using the chatbot for the “standard” scenario can also make the “non-standard” journey self-evident.

This is also a significant prototype validation moment, whereas there are provisions that are too complicated to capture, require too many questions to prompt the user to reveal such “non-standard” conditions, or generally prove the prototype to be non-viable. This is in itself a strong and valuable finding from the prototyping process.

Level 3: Service Journey Map

The third level, building on the previous two, is about diving deep into the context of a certain solution. The students are asked to develop a service map, reflecting their previous chatbot proposal, focusing on the specific steps in the full lifecycle of the proposed service. In the class, we used a service journey map template adapted from This Is Service Design Doing by Mark Stickdorn, Markus Edgar Hormess, Adam Lawrence, and Jakob Schneider and Service Blueprint Template by the Nielsen Norman Group.

This was probably the most powerful prototype as it forced the students to ponder and write down the mechanics of their service, including the human-computer collaboration, sequence, and handover of tasks.

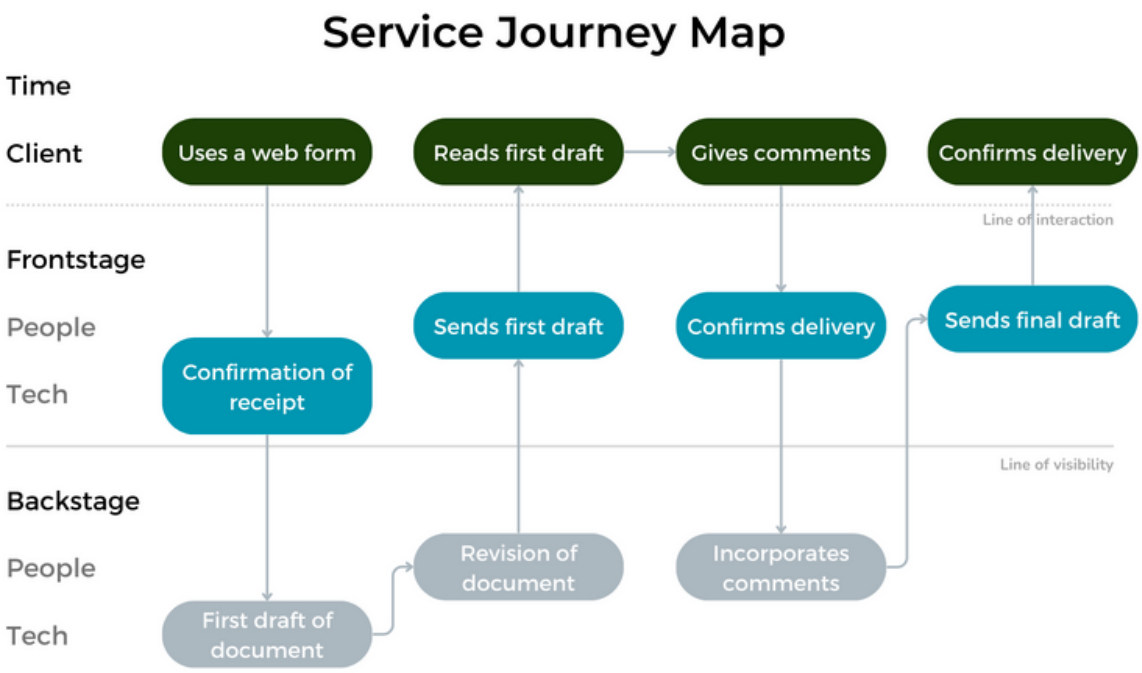


Figure 2: Sample Service Journey Map

Overview of the Process

The starting point of this task is to prime the LLM Application with service design expectations, including the main requirements for the creation of service blueprints. Following up, the students are encouraged to use a short description of their use case to get a first draft of the blueprint. In order to get the best results, the students need to be able to distill the essential rationale behind their ideas.

The first prompt lends itself: “We are developing a [description of the current idea], describe a service journey”. In the case of using LLM Applications with larger context windows, the teams can benefit from the previous priming over the course of preparing the dataset or potentially using it alongside priming on service journey requirements as described above.

Taking the service journey prepared by the LLM Application, the students can then turn it into a visual representation of steps.

Things Line Up Now

In the combination of very clearly defined objectives and a basic outline, the students made much stronger connections between their initial ideas and a realistic service blueprint. This had them explore different possibilities, and also identify the sequence of steps that need to be ideated and designed.

The students specifically appreciated in their feedback that this exercise helped them understand the distinction between the product (the app), the service (the sequence of steps), and the user experience (how we are making the user feel).

This is What We Have and This is What We Don't Have

Based on the first draft service journey, the students also had the opportunity to identify any gaps and dead ends in their proposals. For example, if a chatbot at a certain point acts as a sales funnel that shall at a given moment connect the user with a lawyer, what will happen?

In this respect, the LLM Application can act as a valuable partner in poking holes into the blueprint, resulting in a more holistic model.

Based on the service map, the students can explore different interaction touchpoints. For example, the service map can be transformed into a list of deliverables. From this list, the LLM Application can be then used to build out the specific steps of the prototype and increase its fidelity. That can include generating text for the mockup that will confirm the receipt of the client request, the next interaction between the lawyer and the user, the outcome of their bot, or a list of categories for a knowledge database. Using the Generative AI in this context, the students can focus on the core offering, while also adding an extra realistic touch to their prototype.

Learning Outcomes

Generative AI can be very powerful at the noncommittal prototyping stage, especially when validating early-stage ideas and powering critical review. With limited resources and an abundance of conflicting solution options, rapid prototyping of legal technology can be the pathway to results that best suit a specific organization or user pain point.

Student Impressions

The following is a reflection on the student feedback collected extensively throughout the class. As noted above, due to the iterative design of the class, student experience differed in each semester. The submissions were anonymous and translated from the Czech language.

First of all, the students appreciated the hands-on approach, even when it made the learning experience more difficult at first glance. In particular, multiple students throughout the semesters noted that prototyping, building, and having tangible outcomes from the class gave them a better idea of the practical implications.

> From student feedback: What did you like the most about the course?
>
> The possibility to work creatively, to make something throughout the whole semester. Almost every time I was leaving with the feeling that I created something and therefore I knew that I really did and learned something in the class.

> A broad range of new tools that I can use in practice and the removal of the limitation that I would have to become a legal tech professional to be able to use them

Doing this exploratory work equipped the students with a certain form of creative confidence. With each passing challenge, they would better understand the limitations and strengths of the tools, including LLM Applications, better, and get further in their prototyping efforts. This type of hands-on learning could give us a cue as to how we might be working with legal documents in the next iterations of the tools. In the case of emerging legal minds, this means having the confidence in using the tools constructively, in school and practice.

> From student feedback: What is the most important thing that you have learned?
>
> I learned to look at things differently. When I feel like something is annoying me, I now immediately think about how it could be changed - and most of the time, it is no difficult feat :) The course also gave me a lot of courage to try out a different law-related career.

In particular, the active work with Generative AI resulted in first-hand experience of the balancing act between automation and human judgment, bias and inaccuracy, as well as broader ethical implications. Doing things differently compared to the generated outcomes is an invitation for reflection and an inflection point that needs to be experienced to be fully appreciated. This could be fundamental in empowering lawyers who do not necessarily blindly rely on technology in their practice.

> From student feedback: What is the most important thing that you have learned?
>
> That law can be interesting. That I can find a career even in spaces where it previously seemed hopeless. Accessibility is a much larger topic than I originally thought.

Finally, the class has changed the perception of the students' potential future careers. In the feedback, multiple students have expressed their marvel at the broad array of roles that can be involved in the process of creating innovative legal products and services. The students have highlighted their discoveries that law can be approached and practiced differently. They have appreciated the opportunity to be creative and challenge their previous assumptions about the legal industry.

> From student feedback
>
> It was awesome, nice atmosphere and a unique opportunity to learn something like this. I am super happy that I did this class. It also gave me a nudge to follow this direction when choosing my career.

Implications for Legal Practice

For professionals who already are in the field, these exercises and observations can help enhance the plans for the discovery journey for both themselves as well as for whole organizations.

For example, they could first identify a specific pain point and rapidly prototype how the ideal scenario would work in practice. Just like architects visualize their ideas and make small-scale models, lawyers can benefit from such early-stage rendering of big ideas. By conducting such an exercise, the professional can distinguish between the essential and the nice-to-have, or even test how such a suggestion would be received by their stakeholders.

This could also make for a much more confident negotiator with a legal tech vendor in the future, thanks to a refined understanding of the capabilities as well as a detailed list of requirements for a potential tech solution, derived from the most successful (rapid) prototype.

Finally, these experiments were designed to be accessible with very little resources or previous experience. The tools are widely available and early stage prototypes can be done without using any company-specific or confidential data. This could be an interesting opportunity to introduce Generative AI applications into the organizations, preferably before we reach the point of leveraging large, impactful enterprise-wide technology solutions. By creating this awareness, critical understanding, and confidence early on, legal practitioners can help their teams feel secure and stay focused amidst this great transformation.

Implications for Legal Education

The learning design behind this class was very iterative and experimental and often quite significantly diverging from the more traditional law school format. Below are three key insights that can be adapted by legal educators.

The first thing is the emphasis on agility and iteration in the learning design. In working in a rather rapidly evolving field, no two semesters were the same. The curriculum had to be significantly revised and revisited with each technological advance to stay relevant. To illustrate: at the beginning of the first semester of the Modern Lawyers course, most hadn’t heard of Generative AI. In the second one, most students haven’t used it before. In the third one, students were confidently exploring the limits of GPT. In the fourth, a significant fraction of the students have their own GPT4 subscription and fluently toggle between various LLM Applications. This genesis is important, as every single exploration of these technologies in an academic setting in a controlled environment at any stage can give us more insight into how to adapt the curricula for the needs of future generations of professionals.

Secondly, for legal educators who have been experimenting with or considering including Generative AI in their classes, this is an extraordinary opportunity. The field is, however, leveled, as the younger generations often have the same or even stronger grasp of the technologies as seasoned professionals do. This lends itself to an unusual, collaborative, rhizomatic learning, where we influence each other through discussion and experimentation. It is something to take into consideration while designing the overall atmosphere and tone of the respective courses.

Finally, we cannot truly speak of holistic legal education without proactively addressing uncertainty and ethical considerations that surround the use of technologies in the legal context. The Modern Lawyers course has a built-in approach to talking about ethics, inclusive design, and accessibility. In this respect, the deliberate use of these tools in practice is crucial to the ability to critically evaluate them. Only then can our guidelines and recommendations be informed and help us lead by example over the course of training responsible practitioners who are truly mindful of the broader impact of their work.

Conclusions

As a legal community, we are currently finding ourselves in the discovery stage, where we are among other things defining the most impactful use cases. Using Generative AI for rapid prototyping can give us access to a time-efficient tool for early exploration into real-life applications of the tech in various organizations.

The Modern Lawyers course approached this challenge as a laboratory. The specifically designed tasks in combination with very light software requirements made an accessible way of helping the students to get familiar with the new tech in small, manageable increments.

Any legal use case, be it visualization, chatbot, or a service map needs to ultimately serve a certain need of our clients and users, as well as our communities and society. Generative AI can help us effectively maintain this focus while taking care of the nitty-gritty of early-stage prototyping. Instead, we can focus on testing and discerning the immediate and long-term impact of our solutions. With the prototype in hand, we are much better equipped to ask hard questions regarding our products and services, as well as the future of legal education and training.

Finally, rapid prototyping can contribute to the overall mission of creating well-rounded legal professionals who do not shy away from being proven that they are wrong. Therefore legal educators must embrace this discovery period and create space for their students to undertake experiments in the safety of university settings. With this breath of fresh air, this technology could finally achieve the mindset transformation in the legal industry that we, innovators, have been long dreaming of.

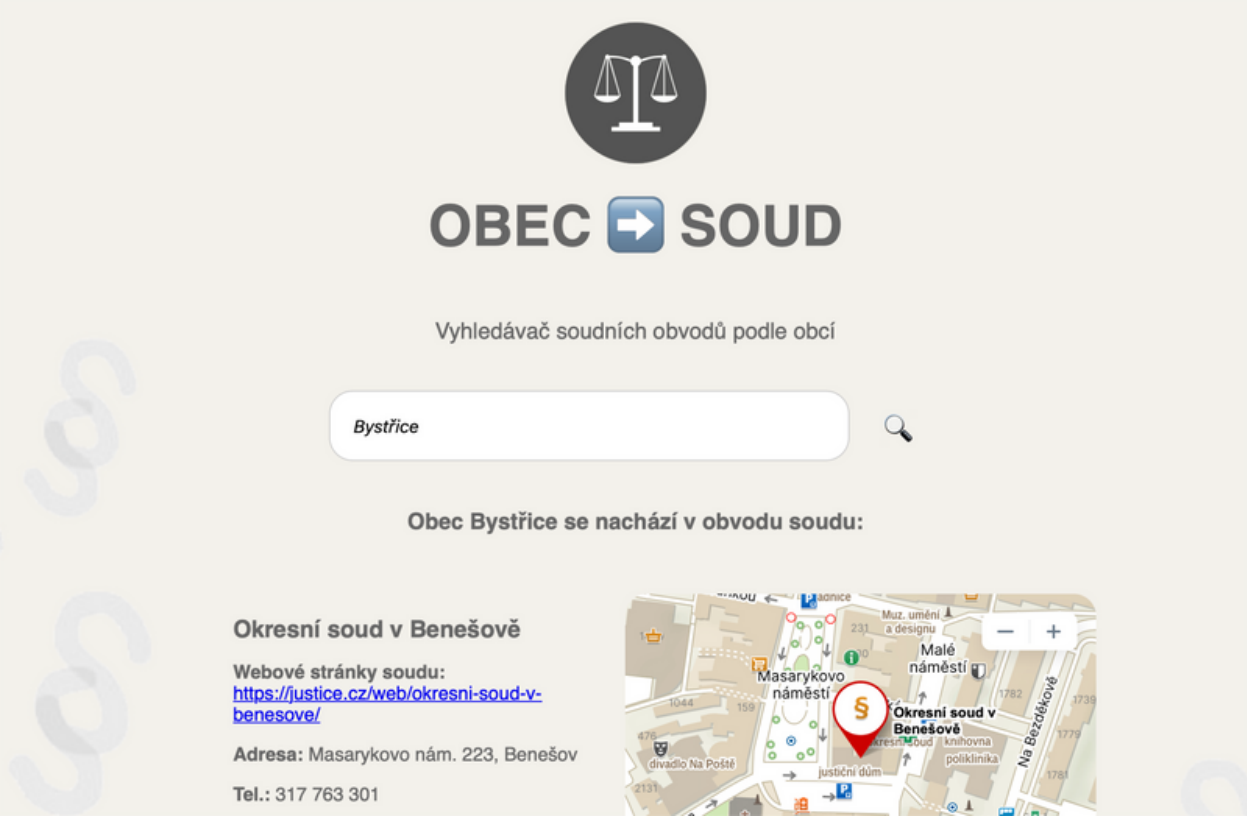
Addendum: Student Examples





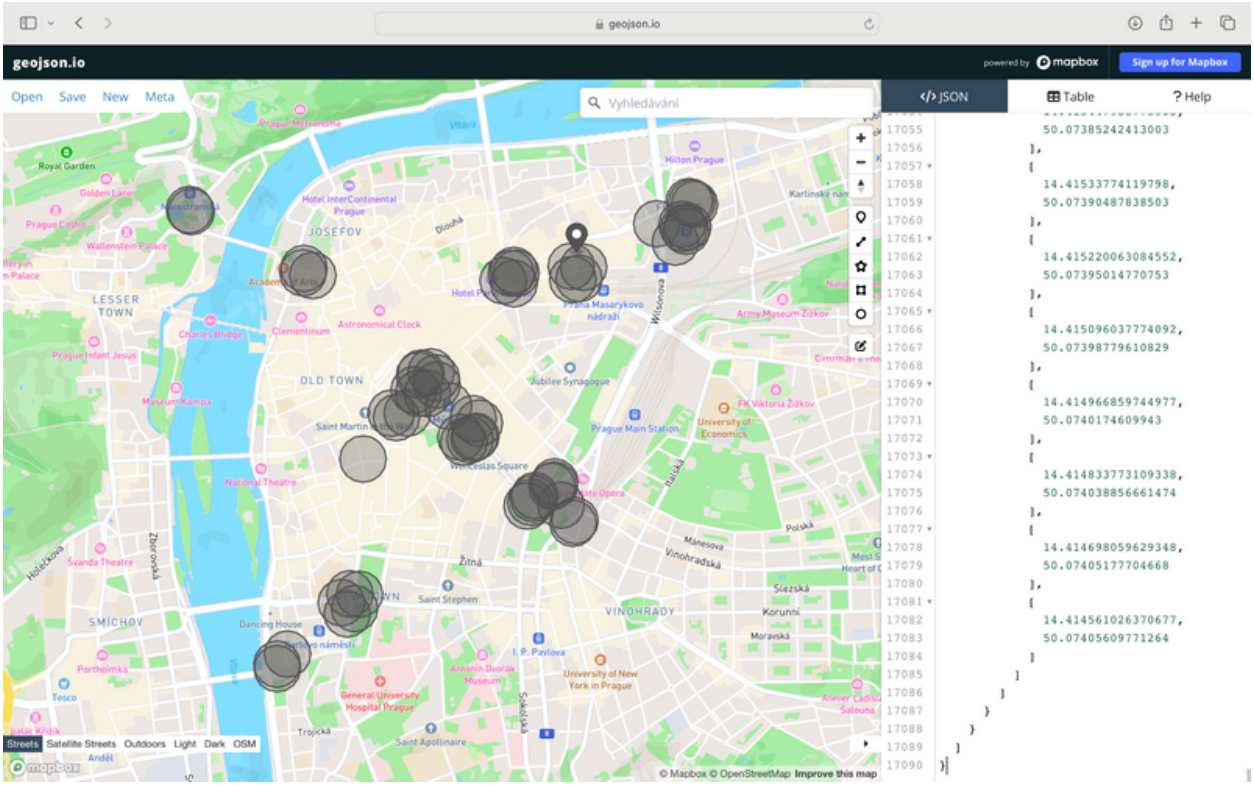


ClearCase prototype by David Bouz, Daniel Řehák, Simona Stasová, and Eliška Syrovátková explored the possibility of managing interactions with a lawyer in an app, including legal queries, case management updates, and budgets. The team was mentored by lawyer, software engineer, and start-up founder Milan Dang.



You can view a video demo of the project here.

Obec -> Soud (in English: municipality to court) is a prototype by Vojtěch Hraběta, David Mužík, Kateřina Novotná, and David Šimek. There are more than six thousand municipalities and 86 district courts in the Czech Republic. The tool helps Czech litigants identify the district court where they need to file their claims, based on the residence of the defendant. The team had no previous experience with coding, and managed to build their prototype using ChatGPT, with a support of a sibling of one of the team members. The team's mentor was Michal Kuk, a legal aid engineer.



The team (Martin Krejčí, Tomáš Moska, and Vojtěch Strouha) took on the challenge of translating a complex municipal ordinance into a user-friendly format. The selected ordinance outlines rules regarding the consumption of alcoholic beverages in public spaces in Prague. This is a regular source of confusion for locals and visitors alike. To address this issue, the team used GPT4 to create an open geoJSON dataset of streets, hospitals, kindergartens, and other noteworthy points in the city center, where consumption of alcohol is prohibited. The team was mentored by data scientist specializing in Prague public data, František Hána.

The author would like to thank the Department of Legal Skills, Faculty of Law, Charles University in Prague, the head of the Department, JUDr. Mgr. Michal Urban Ph.D., as well as the community of students, mentors, guests, and friends, who have contributed to or otherwise support the class.

GPT 4 has been used for a critical review of this article