Global Land Grabbing Dynamics

VIS W16, Written Report M4

Anna Aichinger
Aleksandar Doknic
Reshad Dernjani
Adrian Thöndel

21.01.2017

Table of Contents

1 Introduction

1.1 Project type and project team

The project type is “Design Project”. The focus lies on the visualization design process. The goal is to create the most expressive and easily readable visualizations for a certain problem.

The team consists of four members:

- Anna Aichinger
- Aleksandar Doknic
- Adrian Thöndel
- Reshad Dernjani

1.2 Topic and motivation

"Land Grabbing" is the term for a phenomenon happening globally since around 2000, describing a process whereby private corporations, mainly based in the Global North, purchase vast amounts of land, mainly located in the Global South. The goals are mostly the securing of scarce resources and acre land, in order to yield increased present or future profits. There are often negative effects on local communities, ranging from displacements to environmental degradation.

Because “Land Grabbing” is a process carried out by private actors, there is no official monitoring or tracking of land transactions. Here it becomes necessary to make the published data, however scarce and incomplete, easily accessible for politicians, journalists and development professionals. Our visualisation aims at achieving this goal.

1.3 Problems

1.3.1 Data Set Problem

We found two table format datasets about land grabbing deals. The first one has around 450 rows, the second one over 2000. While the latter has more rows the former has a better description of what the grabbed land is being used for. Sadly, the time data is incomplete in both of them and the amount insufficient for creating time-based predictions.

1.3.2 Interaction Problem

The bundle chart uses a JSON file instead of the CSV file, but the transaction filter merely contains the line in the CSV file as information.

1.4 Users

Everyone who is interested in the topic of Landgrabbing is a potential user of our visualization software, in particular, we expect the following groups to be our users:

- Journalists
- Blogger
- Students

1.5 Tasks

Our visualization is supposed to help us answer questions like:

- Which countries grab land from other countries and which countries are being grabbed by other countries?
- Where are the grabber and grabbed countries geographically located?
- Does [Austria] participate in Landgrabbing?
- Which [german] companies participate in Landgrabbing?
- What is the average buying price per hectare in [Canada]?
- What are the parts of the world that are most affected by Landgrabbing?
- What are the parts of the world that participate in Landgrabbing the most?
- Which country grabs the most land in hectare?
- Which country is being grabbed the most?
- Which continent does [China] have the most deals with?

1.6 Data

In order to create a visualization of the Landgrabbing phenomenon we need information about the deals between involved countries. Besides the deal data we also need general country information to properly address the user tasks that involve any information about the country besides the name and location. The data required to assign the country names to a geographic location can be delivered by the library used for visualizing world maps.

2.3.1 Landgrabbing Data

Source Option 1: https://www.grain.org/article/entries/4479-grain-releases-data-set-with-over-400-global-land-grabs

The data from grain.org could be used as the data source for our visualization project.

Pros:

- shows the grabbing countries

- shows the involved grabbing companies and sectors
- shows the amount of hectares that has been bought
- shows the status of the deal

Cons:
- Data is from 2012
- Data only includes deals from 2006
- Does not show year of the deal in own column
- Does not provide much information about the projected investments
- Only XLS, so it has to be converted into CSV and then uploaded it into the GIT Repository.

Source Option 2
http://landmatrix.org/en/get-the-detail/

Pros:
- shows the grabbing countries
- shows the exact location
- shows over 2400 deals, compared to ~400 in the dataset above
- shows status of the deal
- recent data up to 2016

Cons:
- intended size and contract size are separated

Even though we initially started with the Source Option 1 in mind, we decided to switch to Source Option 2 as it would significantly improve the expressiveness and qualityof our visualisation.

National Data

Source: http://stats.oecd.org/

We can connect the involved countries with official OECD statistics in order to view information about the selected countries economy.

Pros:
- A huge variety of statistical data
- Highly reputable and reliable source

Cons:
- None

2 Related work

2.1 Landmatrix.org

The website http://landmatrix.org/en/ offers several landgrabbing visualization. They use a different dataset than we use.

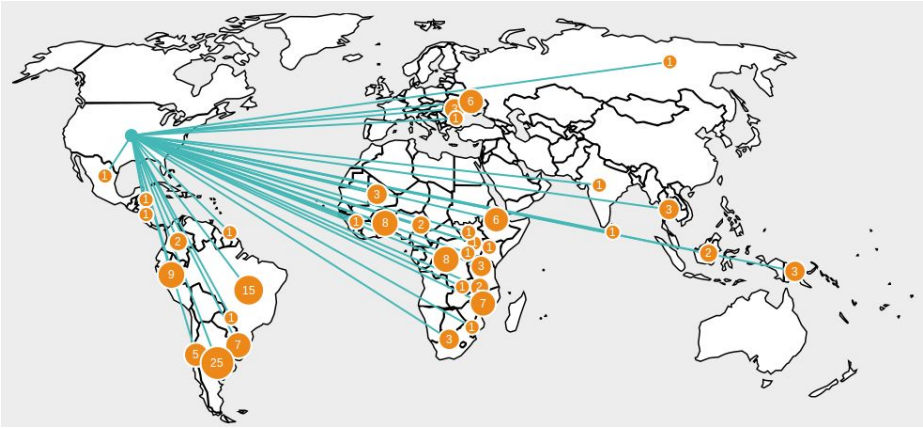

The “global map of investments” relies on animation in order to view the numbers of transactions associated with a certain country for a short time. It is similar to some approaches that we considered for our visualization. One of the disadvantages of this visualization is that the mouseover elements are sometimes difficult to associate with the country beneath. They also overlap, particularly at the beginning of the animation. The whole visualization is detailed but not as readable as it could be.

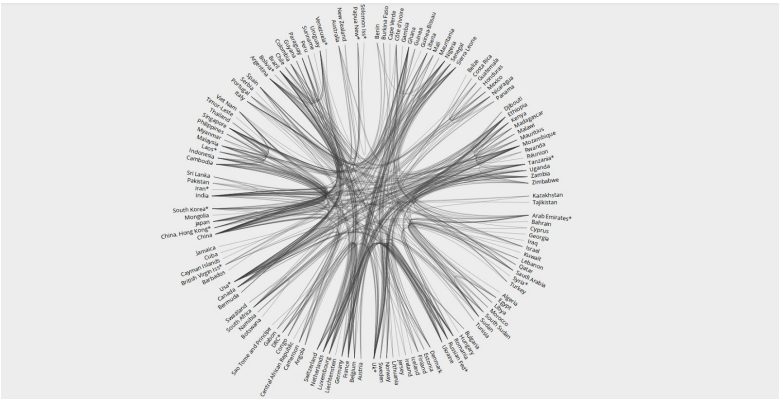

The “web of transnational deals” has similarities with our bundle chart/transaction view.
It is generally a good view, but it does not give the viewer a feeling for the geographic location of the countries. In our visualization each of the countries is linked to the map view.

It should be noted that our visualization was designed before we were able to compare it to landmatrix.org, therefore it is very interesting that we invented similar visualization approaches. We put more focus on an integration between the both aspects (countries and transactions) of the data than the landmatrix.org visualizations.

2.1 petergiovanni.com/Landgrabbing

This website only has a static cartogram from countries that were grabbed from 2006.

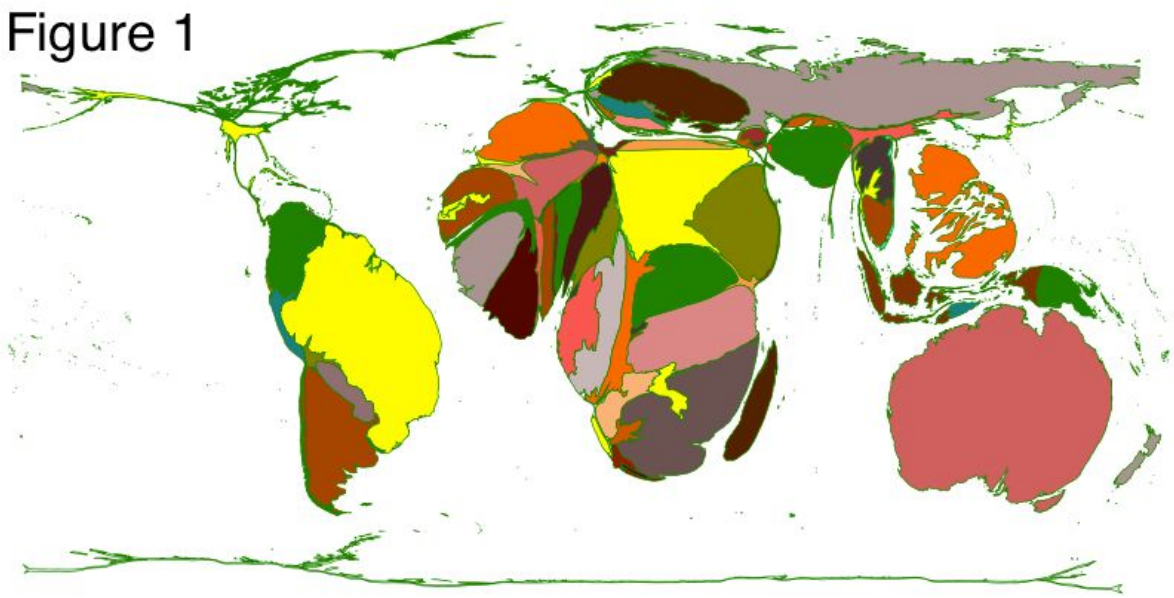

Figure 1. Distorted world map according to the relative amount of land grabs in each country. In other words, the area of each country was rescaled in proportion to the amount of land that has been acquired by foreign investors after 2006. The colors in the cartogram are only used to distinguish different countries.

It only shows the grabbed countries and not the grabber countries but it is still very expressive, although it can be difficult to figure out which country is which.

3 Approach

3.1 Considerations

Since we have a design project with specific user tasks that our visualization needs to address by using complex data we will need more than one view. The sum of all view elements linked together is referred to as “Dashboard”.

Our visualization has to include the following important aspects:

Countries

- Geographic location
- Country as a buyer
- Country as a seller
- Relation to other countries
- Economic data

Transactions

- Buyer country (Grabber)
- Buyer company
- Seller country (Grabbed)
- Transaction size financially
- Transaction size in land size (hectare)
- Usage of land

We do not include certain aspects in our visualization as they are out of the scope of our project. This includes, for example, implementation status, negotiation status and the nature of the deal. Our visualization only shows the concluded deals.

The transactions could be viewed as relations between two countries. A relation can be depicted by drawing a line between the participants. Therefore a country can be viewed as a complex type of node, while a transaction can be viewed by as a complex link. Both contain more information than a simple node-link graph could show at glance, however, it is essentially a link-node problem that we are facing. It is not a simple network but it can be considered a graph with many many independent 2-node trees and leads us to the same problems.

We want to show both aspects of the data: The country information and the transaction information. The easiest (and therefore the best) solution would be a unified visualization that shows us both aspects.

Using a World Map means showing countries (the nodes) bound to their geographic location on a map and using color channels to show some information about each country. That would be the most obvious and easy-to-understand solution when it comes to visualizing countries.

However, adding the transactions, hundreds or thousands noods across the map, no matter what kind of map is being used, will inevitably make the view unrecognizable (hairball problem).

The scale is approximately 200 nodes (more if we include the exact location) and approximately 2400 links. Even using a diagram with free spatial channels we would face a very complex network. We could filter away some links to make it more readable but then we would lose the option of showing a complete view of the situation at once.

A Matrix View would be a way to visualize thousands of transactions, links at once without creating a hairball. We would end up with an approximately 200-dimensional square a matrix. It would be difficult to sort the categorical elements on the axes in any useful way.

Using a Hierarchical Bundle Chart would also be a possible solution for visualizing a big number of nodes and links in a structured way. We could show the flows of land from country to country, but the countries would be mere labels without geographical location.

Our dilemma as the visualization team is now to decide between showing the countries with all their underlying relevant information, and showing the transactions with all their underlying relevant data.

In general, a Hierarchical Bundle Chart would be sufficient to solve most of our tasks, but it would fail to create a sense of location and make the whole visualization, which is about a global topic, more abstract and difficult to imagine. It would also be unable to address user task 3.

Using both charts seems to be the only way that allows us to show a complete view both aspects, and even then we would need mouseovers or popups attached to the countries and transactions to show a deeper level of information.

An alternative approach would be making the world map a passive chart without any interaction, showing only the countries that are currently involved in the transaction bundle that is currently selected with a mouseover on the Hierarchical bundle chart. That way the focus could stay on the Hierarchical Bundle chart, but we could relate the countries to their geographical location.

The default view of the world map would be showing the grabber countries by color. The specific view would show a selected country (white) and the countries where they grab land or from which land is grabbed from them coded by color.

3.2 World map

A world map can be depicted in many different ways. While in our case the Cartogram might be an interesting way to depict country sizes, because it'd automatically only use the space for relevant countries, we would have a problem when trying to compare different data sets with each other, for example, grabber countries and grabbed countries, or particular countries that are involved in grabbing, because the map would look differently each time.

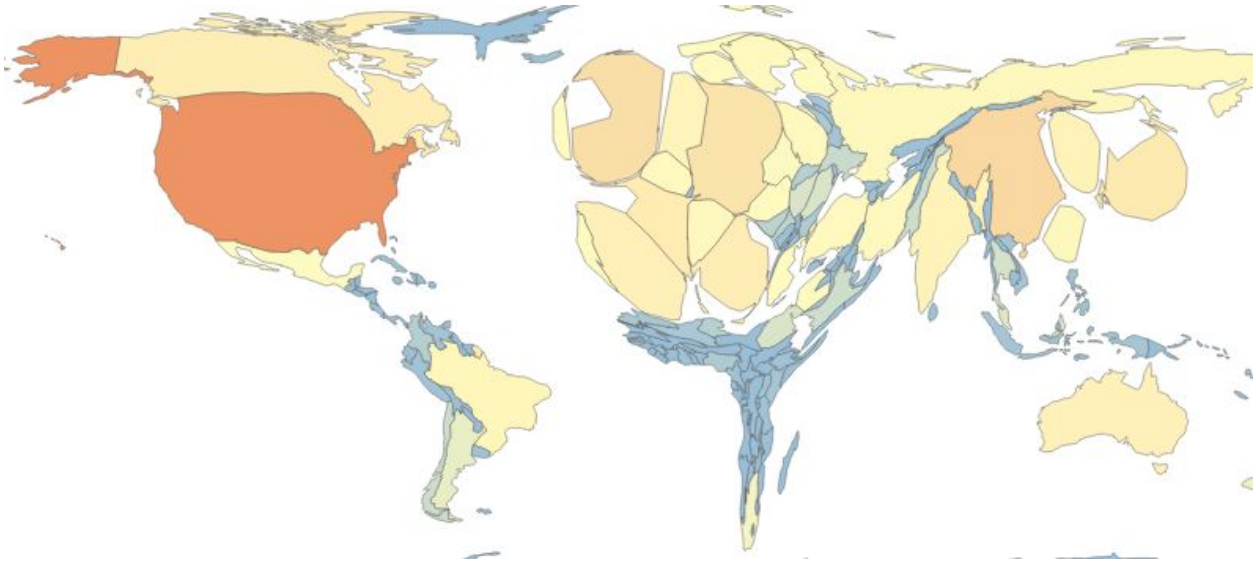

Mercator Map might not be the best choice when visualizing something that involves country areas.

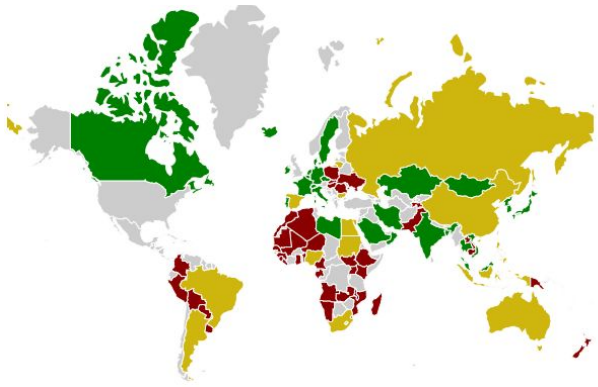

The basic map is a more proportional map based on mercator map.

Mini Maps would be the best choice to offer a side-by-side comparison instead of showing a colorful mess.

Our choice depends on the Dashboard that we decide to use, but it seems that proportional mini maps would the best option for a “showing” map without interaction.
The hue of the color channel is used to encode whether a country is a grabber or a grabbed country while the luminance is being used to encode the intensity (by hectares) of grabbing or being grabbed.

3.3 Country Info Box

Each country will show specific information on GDI (Gross Domestic Product) and HDI (human development index). The data source for the country box will be as mentioned: http://stats.oecd.org.
On mouseover the tooltip of the specific country will be displayed these information. We will use the a tooltip for this task because many independent information are displayed which can’t be useful implemented in form of a bar chart, histogram or pie chart.

Country: Germany
GDP: 3.352 tn. US$
HDI: 0.916

Tooltip country info box

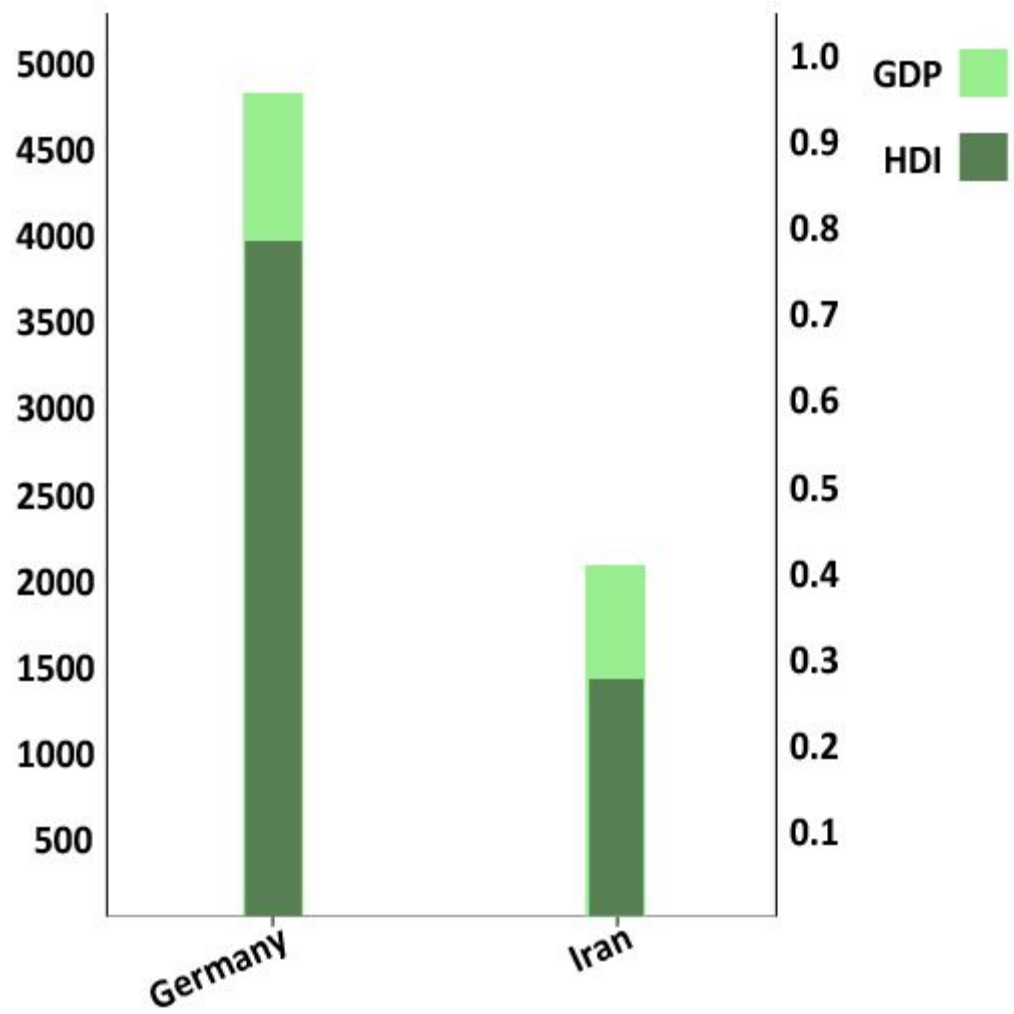

Version 1: Stacked Bar Chart County Info Comparison

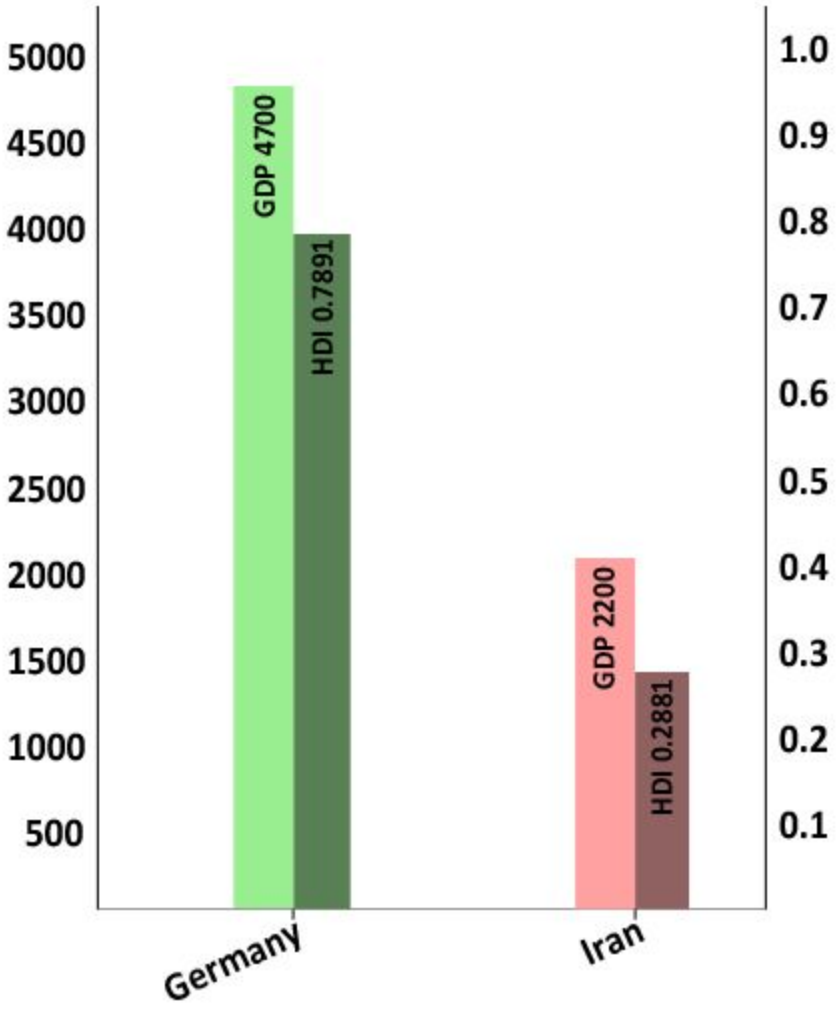

Version 2: Bar Chart Country Info Comparison

The country info can be displayed even when countries are selected from other interactive actions on this visualisation e.g. selecting on the bundle chart one path. Than there are two versions of bar charts we could use. A histogram or pie chart does not make sense at all.
Javascript in combination satisfies us with functions to create the tooltip and the chart.
The stacked bar chart is not easy to read. It is hard to differentiate the GDP and HDI of each country. The best solution here is the normal bar chart.

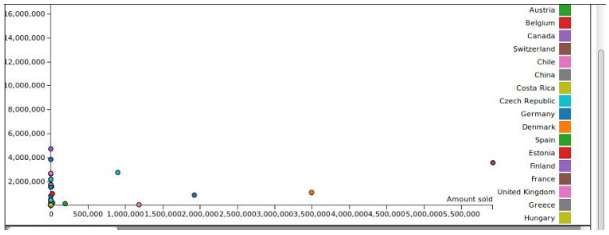

Version 3 (Final): Country Info Comparison in a Scatterplot

For the final implementation we had to merge those options mentioned in a Scatterplot. Here each country is represented with his gross domestic product and the amount of total sold land in hectare. By mouse over you will see the values in detail. It is possible to compare all countries how the amount of sold land is correlating with the gross domestic product. Therefore selecting a country in our bundle chart will filter only the countries correlating in terms of transaction, there it eventually is possible to see what gross domestic product has the seller or buyer of land.

3.4 Transaction Info Box

1) Which user tasks are addressed with this visualization?

- How much land (in hectares) is being grabbed? (in absolute / relative terms)

By selecting a country or a specific transaction, this view yields either the grabbed hectares in total or the respective amount of the transaction.

- How much money is being spent on the land?

By selecting a country or a specific transaction, this view yields either the money spent for the hectares in total or the respective balance for the hectares of the transaction.

- Which economic sectors are involved in a deal, and to what degree?

  Users can query this information by selecting either a country or an adhering transaction.

2) What kind of data is used in this visualization?

Since the transaction box shows detailed information about a specific deal, a predefined set of data will be visualized out of the source. This includes information about how much money was spent in the selected transaction, what are the involved countries, which company signed the contract, how many hectares have been grabbed, for what purpose is the acquired land going to be used and to which economic sector does the transaction adhere?

3) What charts are are used?

Throughout the elaboration phase there have been experiments to visualize data with different chart types where at first bar charts and histograms have been taken into account.

We first tried to apply some information about production purposes to histogram charts to indicate which goods are produced the most by one company. However, companies may have various purposes, why the bin-value for the histogram has to be re-validated for every visualization and if not done so, important data may be cropped off to the wrong bin. For this reason to visualize this kind of information a histogram won't work in our project.

Since the information we are going to use in our project has many aspects that can be subdivided in categories and the resource has limited occurrences of e.g. economic sectors

which can be summed up into categories respectively, we decided to use bar charts instead, because they work for categorical data excellently. A single transaction possibly can adhere to any of the defined categories but only covers a part of the whole data set.
In order to provide different insights to the data the transaction view will contain different versions of bar charts that display information in various ways.
Another advantage of bars is that if charts are drawn with the same dimensions, they can be easily compared by users. Additionally, source data can be combined with new data in every chart.
This condition will be leveraged in the detail-view of one transaction by using two bar charts. The first indicates on which products the company has specialized by combining all transactions of this company. The chart next to it sums the data of all transactions of the selected country and faces it to the count according to the deals a company has made in a specific sector.
This partitioning grants insight to the company's production purposes while providing an overview on how many deals this company has made in particular and how active this company is compared to all other grabbing corporations of one country.
However, if the visible data does not have much aspects in common or can not be applied to the context of the source data, the use of two bar charts is pointless here.
We decided to develop the transaction box as a tool for comparison of a company's activity with the grabbing behavior of the whole country. This would be a good idea to indicate how big of a deal this transaction was for either the country or the corporation.
To achieve this, pie charts are used. Well, an often experienced disadvantage of this radial chart type is that accurate perception of the layout is a problem. However, in this scenario our pie charts are neither intended to give detailed information nor be used for user interaction but rather to give a quick overview of the transaction's extent according to the grabbed hectares. The whole pie represents 100% of all grabbed hectares, whereas a slice implies the amount of hectares for one transaction. A pie chart does a good job here, to quickly show how big of a deal this transaction was for the company or the country itself.

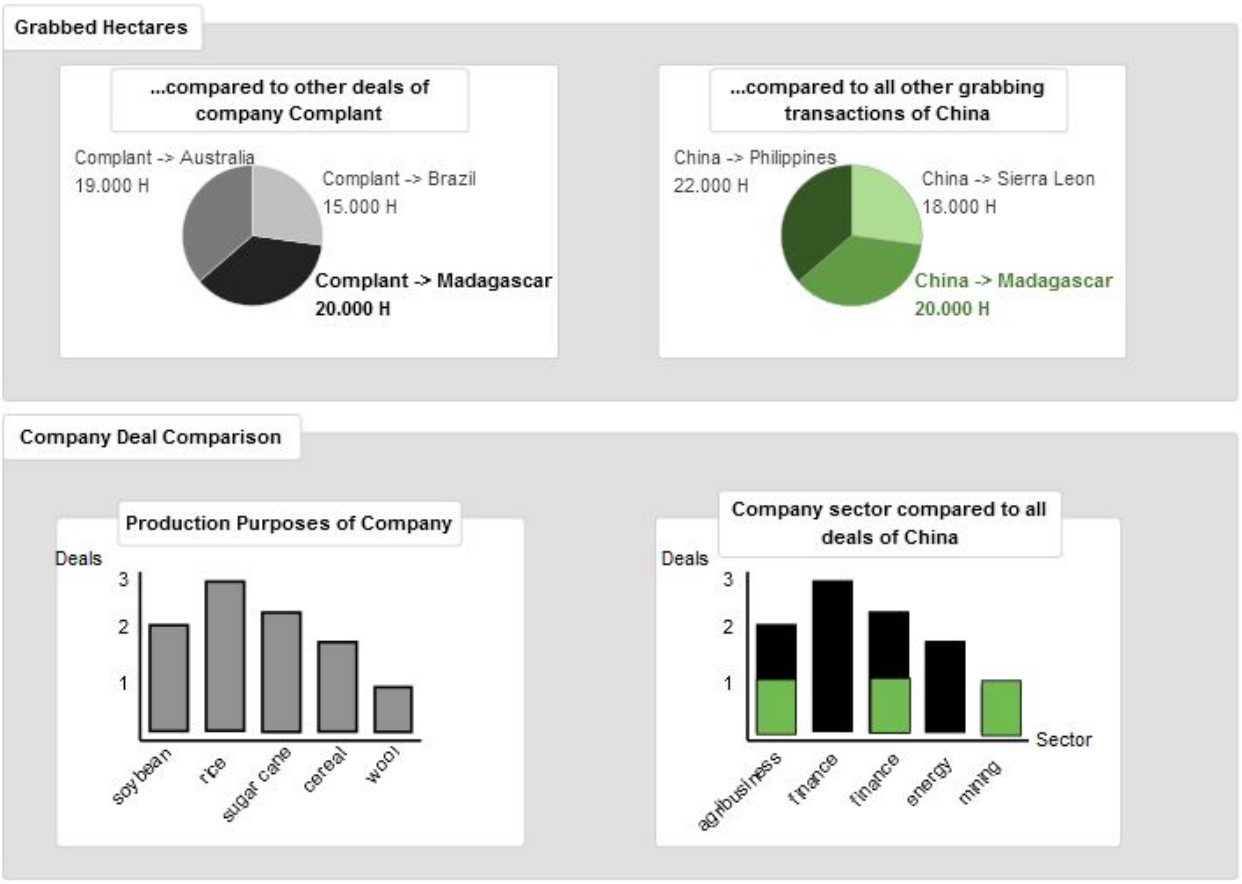

Like this figure implies, the transaction-info-box visualizes data with pie- and bar charts.

4) How can user interaction be performed?

As soon as a user has performed an action on either the edge-bundling chart or the world map (depends on the dashboard) and a country is selected, the transaction-info-box pops up and the required data is brushed to the info box. Data is shown and every deal that adheres to this country is listed in this box. Users can click on one transaction to request more detailed information about this transaction.

3.5 Visualization Design

The fundamental problem of our design study is the visualization of two different concepts that are relevant for our user tasks, countries and transactions.

During the M2 presentation feedback the Flow Map, as described by papers like Edge Bundling in Information Visualization by Hong Zhou, Panpan Xu, Xiaoru Yuan, and Huamin Qu, was proposed as an alternative approach to our separated views for countries and transactions.

It is true that the edge bundling applied on geographical maps brings structure into a number of graphs that would otherwise end up as a visual hairball. It is, for example, an excellent choice for the visualization of migration patterns to a particular place, as illustrated by page 5 of said

paper. The usefulness for visualizing a number of mostly independent flows with multiple starting points as they occur in our project remains to be proven yet.
However, assuming that the application of edge bundling solves the hairball problem on the integrated visualization and enables an integrated view that shows the transaction edges in a structured way on a world map, we made the following pro/con considerations for our old and separated approach, and the new and integrated approach:
Integrated Country-Transaction Map

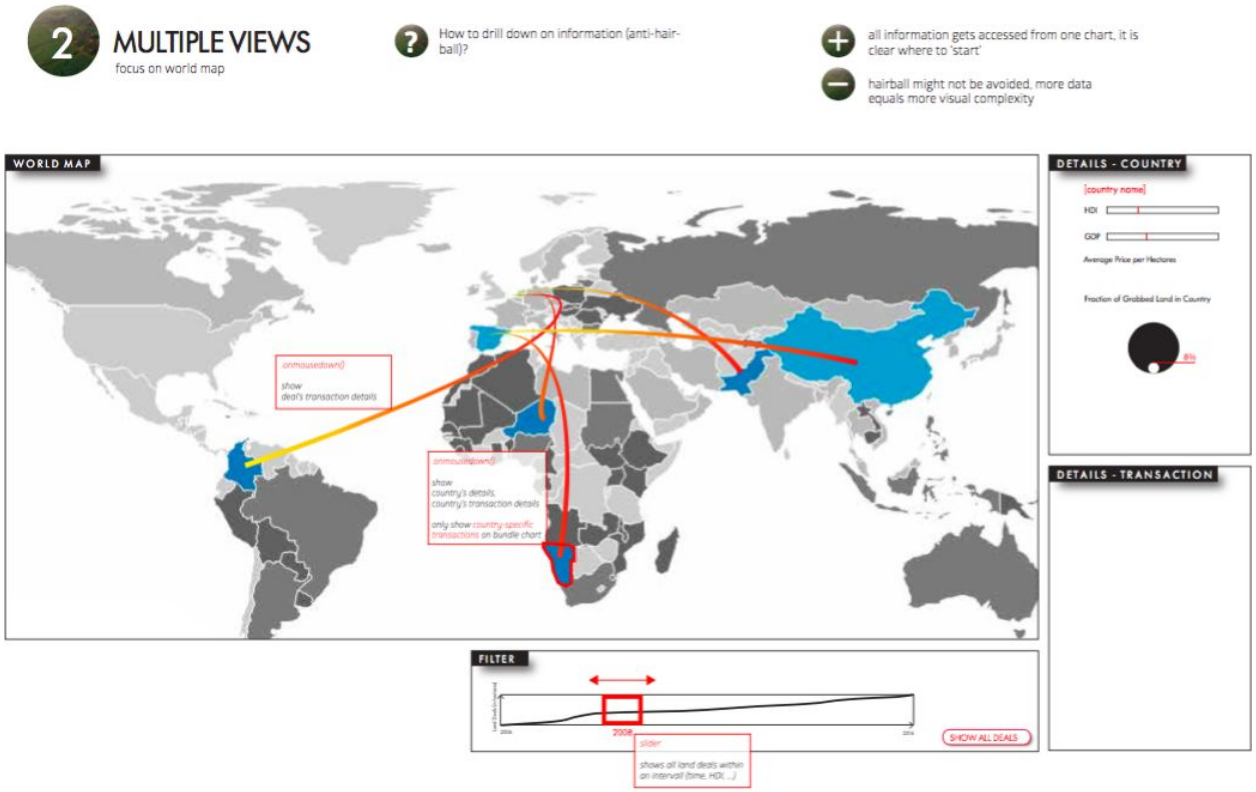

In this approach, as considered in M2 Dashboard #2 and proposed in the feedback, we use a single view that shows a world map with color-encoded countries as well as transaction lines that connect all grabber countries with their respective grabbed countries for each and every transaction. Those lines are supposed to be organized using edge bundling.

- Everything is in one view, less work for the eyes.
- The geographic relation between transactions and geographic location becomes clear at one glance. This advantage only holds true for the whole unfiltered view, because selecting a particular country shows the transaction partners even without edges.
- Using color encoding (i.e.) to show the grabbing intensity or other properties of a country on the map becomes difficult because of the mixed countries that do both, buy land and sell land. The luminance can be used to show the intensity of grabbing done by a grabber country, and grabbing done to a grabbed country, but the aggregation of both for a mixed country would not make much sense.

- Edges are (potentially curved) lines and therefore have to start from a particular 0-dimensional point in a country and end in a 0-dimensional point in another country. In order for the edges to contain any buyer/seller information they should be color-encoded. The problem of multiple overlapping edges of the same type can be solved by using algorithms to aggregate them into a single edge with more width, but the problem of overlapping edges of opposite types with opposite encodings are more difficult to solve in a simple and visually reasonable way.
- A rather minor but notable issue is that the edges would either go across the country areas or across the borders, potentially covering/overlapping one or the other which might reduce the expressiveness of the visualization for smaller countries.
- Properly adapting, optimizing and applying the map edge bundling algorithms from various papers and tutorials to our current visualization task might require non-trivial algorithmic and technical considerations and focus on implementation rather than design.

Separated Country-Transaction Views

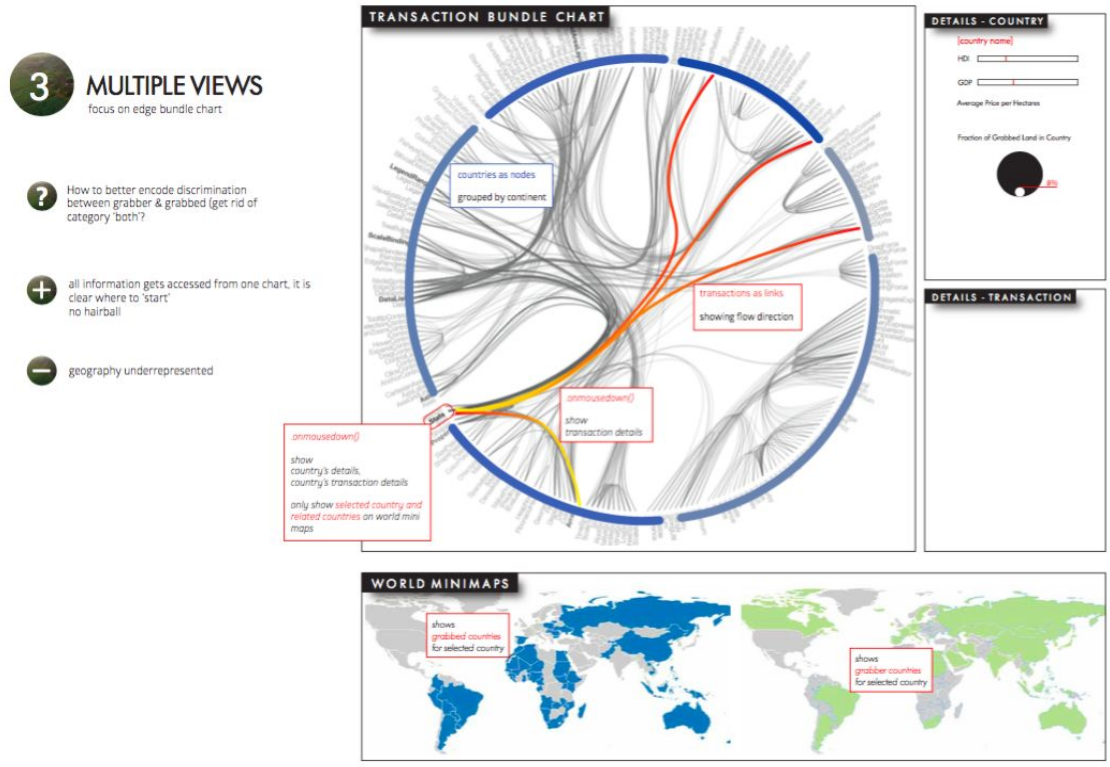

In this approach, as described in M2 Dashboard #3, we use a Hierarchical Bundle Chart to depict the transaction flows using a circular abstract arrangement of countries, and additionally one or multiple separate world maps to support our geographic tasks and offer a general view on the grabber/grabbed countries as well as a geographic view for particular countries and transactions. These world maps and the HB Chart are tightly linked.

- KISS principle: The view uses well known, reliable and well-established visualization elements where each view has a particular task.

- Using separated views enables us to avoid sacrificing encodings like color or space just to maintain the visual integrity and readability as it is required by more complex charts.
- One powerful aspect is the dynamics of structuring/ordering the data and showing flow patterns. The Hierarchical Bundle Chart countries can be sorted by continent, GDP high-low, GDP per capita, most-grabbed-least-grabbed, biggest-grabber-smallest-grabber, etc. This would be impossible with the integrated view since the country nodes would be fixed.
- The possibility to use multiple small world maps would not only make grabber and grabbed visible at one glance but also enable the usage of unambiguous color encodings with a clear meaning, depending on whether it's shown on the grabber map or grabbed map.
- Easy to implement (löschen bei Abgabe)
- Easy to separate work tasks (löschen bei Abgabe)
- The relation between transactions and geographic locations is not visible at one glance.
- Requires more space and leads to smaller views.

The first and strongest positive point for the integrated view can be related to a design principle that was mentioned in the lecture of Nov 22th. In the PDF Facet Into Multiple Views on page 5 it is described and emphasized with bold letters that the principle Eyes Over View is important when combining views and that "two simultaneous views have lower cognitive load than remembering previous view".
In this case the advantage of "less work for eyes" is a minor one. On the other hand, many of the issues of the integrated view would have to be solved by offering a "switch" button and requiring the user to switch, for example, between the grabbed map and the grabber map in order to, for example, solve the visualization issues of mixed countries and overlapping opposite transactions.
The last and strongest negative point for the separated view is that it "requires more space and leads to smaller views". The same page in the same PDF describes this as the real-estate trade-off. If we prefer eyes over memory then we have to inevitably sacrifice space.
The 2nd positive point of the integrated view and the first negative point of the separated views are complimentary, it is the relation between transaction and geographic location. Therefore we end up with two questions:

1. Do we even have enough space to visualize the separate views? One major advantage of such an unified/combined view is saving space.

2. Is the aspect of relating transactions and geographic locations, which is the major advantage of the integrated view, important enough to outweigh the benefits of the separated view, the visual and technical simplicity, readability, eyes over memory appraoch, ordering of country nodes in the hierarchical bundle chart, and more powerful color encodings in each view?

The answer to the first question is: Yes, even a smaller maps and bundle charts are expressive and readable and support the user tasks. We don't have to necessarily save space, especially not on our high priority views.
The answer to the second question: It depends on the user tasks. Our user tasks focus on the location on the grabber countries and grabbed countries. We can show this perfectly fine without showing the flows geographically. Therefore the answer is no, it is not crucial to our visualization.
It would be a very pretty and impressive visualization for sure, but for this particular visualization the benefits of an integrated country-transaction view offer only a limited meaningfulness and usefulness which does outweigh the severe drawbacks in the visualization design.
Scatterplot
Another feedback that we received recommended the usage of scatterplots in order to show trends and make predictions.
The most interesting thing to show would be how certain Landgrabbing events change with time.
Aspects:

- Time (available?)
- GDP total
- GDP per capita
- HDI
- Distance to equator
- Average price per hectare
- Amount of land grabbed
- Amount of land sold
- Country area

Our goal is to show the relations between different aspects of Landgrabbing in general and in relation to particular countries. The aspect of time, if available, could be used to make predictions.

GDP - Amount of land bought
GDP - Amount of land sold
HDI - Amount of land bought
HDI - Amount of land sold
GDP - Distance to equator
Distance to equator - Average price
There could be a combox for the x-Axis categories that allows, and a combobox for the y-Axis that allows the user to choose which relations are show. Alternatively, a fixed set of combinations can be selected.

4 Implementation

4.1 Technologies

There are many visualization tools on the market. The only question we care about is whether those tools fulfil our requirements: Free, Easy, Powerful

Tool	Free	Difficulty	Powerful	Suitable
Tableau	Community	Easy	No	It offers most of the view elements that we need to implement, but the dashboard it is not flexible enough for our needs.
Pure JS	Yes	Very High	Yes	It can be used for almost anything, but only makes sense for very specific tasks that have not yet been implemented in any other way.
JS+D3	Yes	Medium	Yes	Due to its popularity and the crazy number of tutorials it has become a very powerful library that is capable of providing anything we would need for our visualization.
vis.js	Yes	Easy	No	A JS library for visualizing data, similarly to d3. However, it seems to be focused on simple graphs and networks and does not seem to provide geographic maps which would make the

				implementation more complicated. A combination of D3 and vis.js would be redundant.
DataMaps	Yes	Easy	Only Maps	A d3-based library that creates geographic maps. Very convenient for its field of usage.

4.2 View 1

4.2.1 General design

The hierarchical edge bundle chart is intended as the base chart of the dashboard. All other views react to the user's selection in the edge bundle chart.

All countries in the dataset (grabbers and grabbed) form the nodes of the charts, the links signify land transaction deals between countries. The links are directional, meaning that two colours are used to signify if the deal is 'outgoing' (land is being grabbed from a particular country) or 'incoming' (land is being grabbed by a particular country).

In this way, each land transaction deal is clearly shown, while the edge bundle chart would allow to see patterns within the dataset easily and at one glance.

4.2.2 Data

The view does not need data from any other views, as all other views are dependent on it. For the creation of the node-to-node relations, a .json data structure was chosen. In the data file, each transaction is equal to a JavaScript object, with its fields containing all the relevant information concerning a deal.

4.2.3 Implementation

The chart was implemented using d3.js only. The nodes are positioned before the links are bundled, using the cluster layout. The layout is then applied to an array of links, each of which has a source and target property. These properties point to the nodes at the two ends of the link. The two nodes each need a parent attribute that points to the parent node in the hierarchy. This is how a 'direction' is created for each land transaction deal.

4.2.4 Default View

By default, the edge bundle chart shows all nodes and links in “neutral colours” :

- nodes in grey
- links in blue

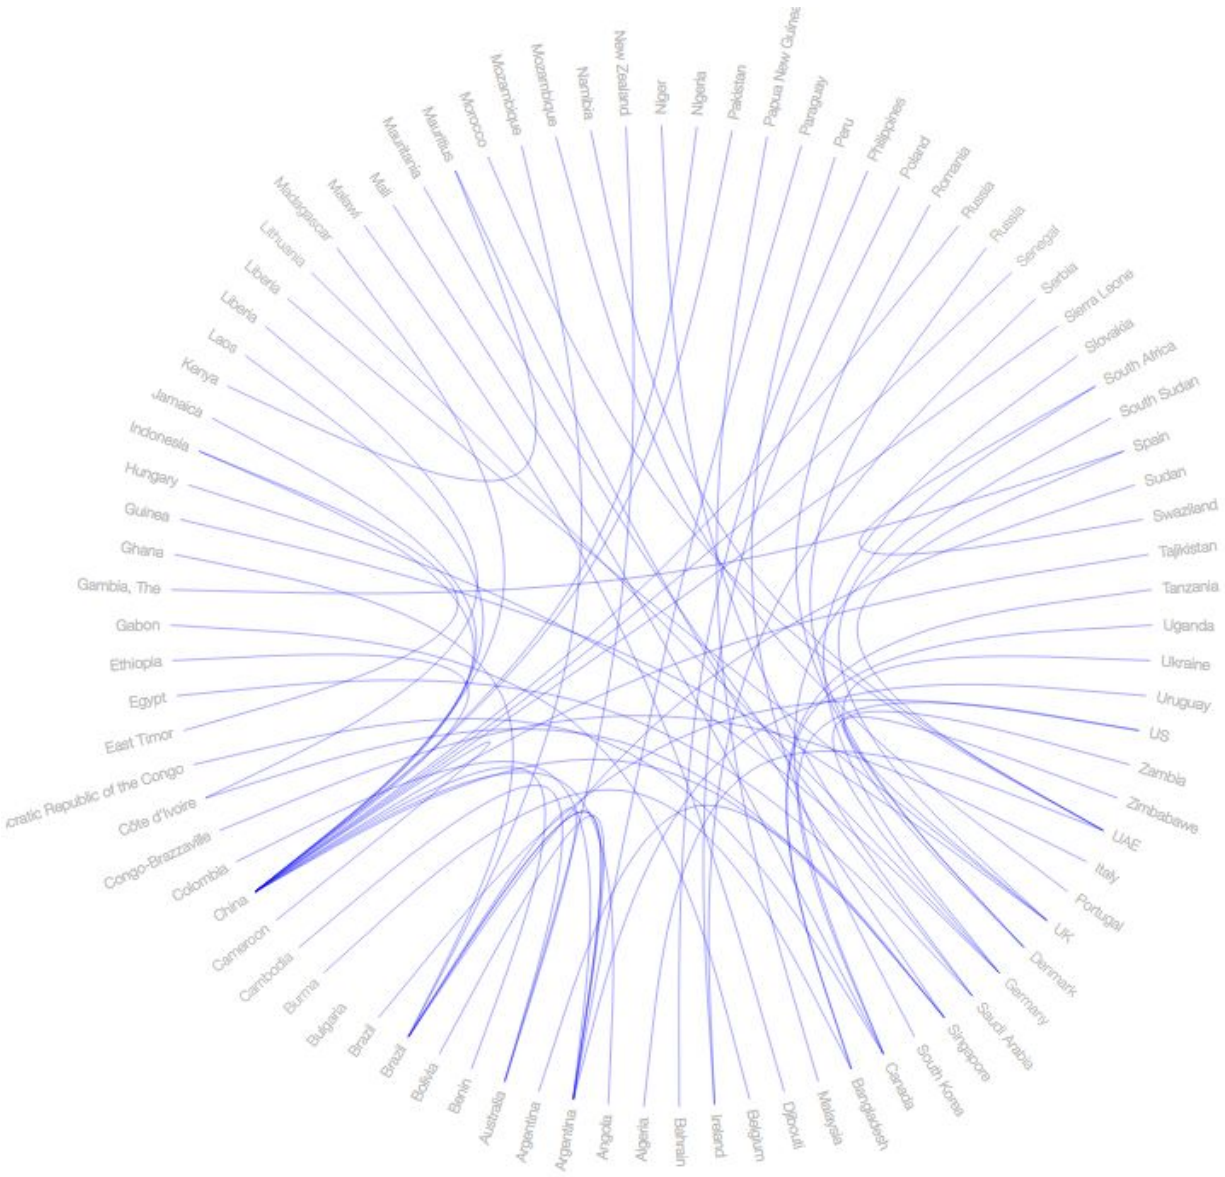

4.2.5 Interactivity, active

When moving the mouse over a specific node, the node and all incoming and outgoing links are highlighted. When clicking on a specific node, country specific information is shown in the transaction info box (view 3) and scatterplot (view 4). When clicking on a transaction link, transaction specific information is shown in the transaction info box.

The colours of links are chosen according to the ‘grabbing’ direction. Links highlighted in

- red signify that the selected country grabs in all linked countries;
- yellow signify that the selected country is being grabbed by a linked country.

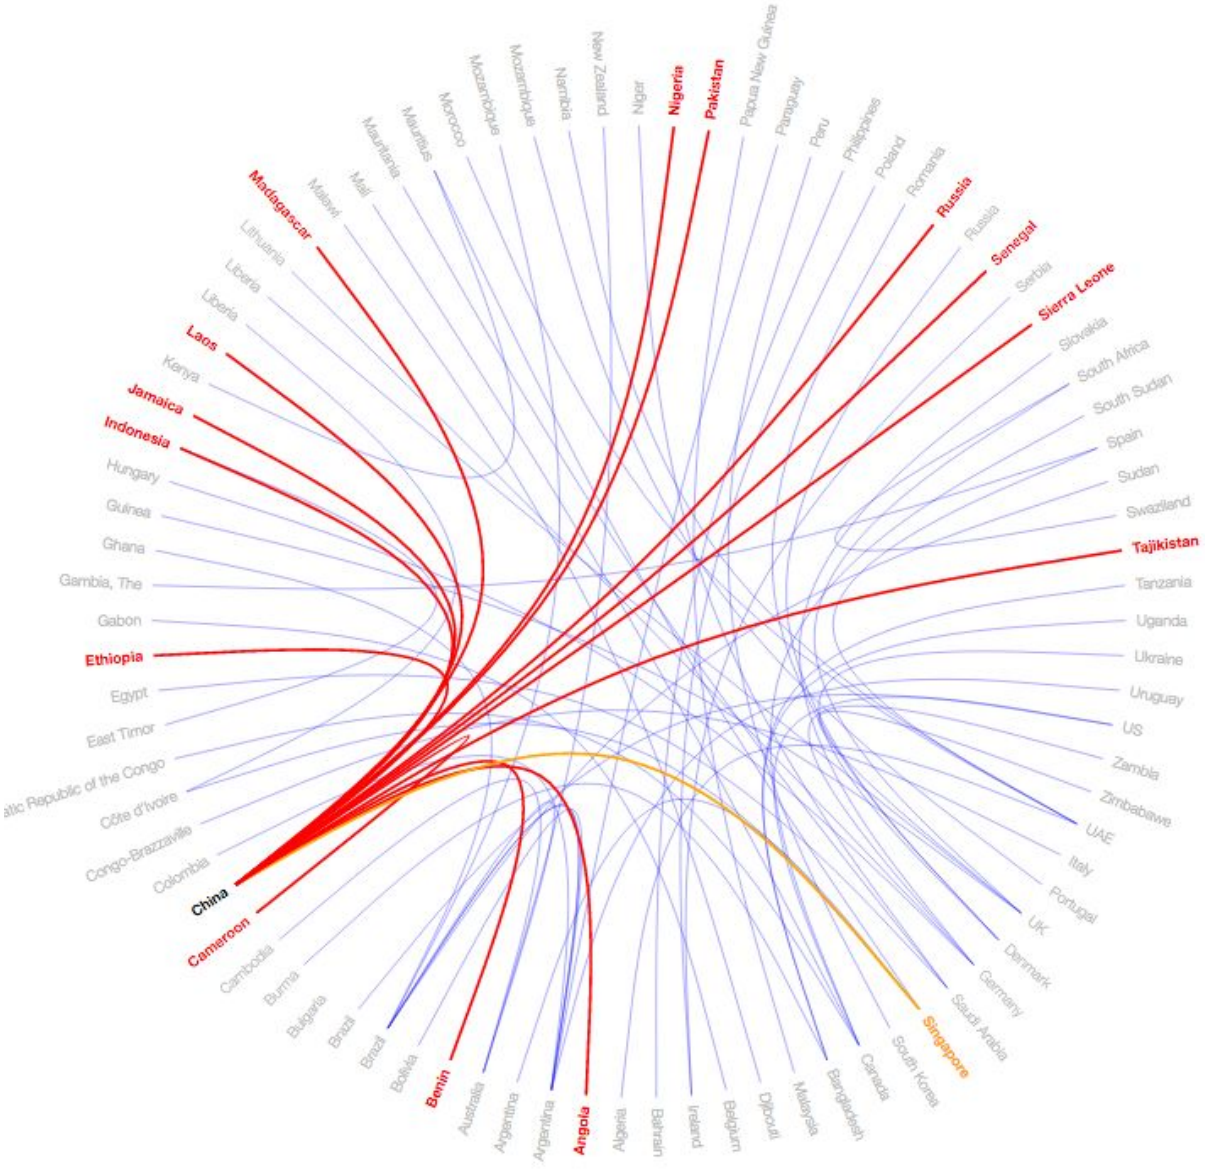

4.2.6 Interactivity, passive

If a country is selected in any of the other views, the country and its incoming and outgoing links are highlighted in the edge bundle chart (see above).

4.2.7 Challenges

The main challenge was figuring out a way to make the data fit the node-link relations required in the code. Because every link should be clickable and linked to its related information in other views, and should be provided with various linking and hierarchical properties (parent/child/source/target), each transaction was stored as a unique JavaScript object in the dataset. Each JavaScript object contains a field for the grabbed country (name), and one for the grabber country (imports).
Also, the dataset contains a couple of discrepancies in spelling and naming of countries. This was solved using a helper script that assigned fixed string values to each name field.

4.3 View 2

View 2 shows two world maps next to each other.
The left map shows the grabber/buyer/investing countries.
The right map shows the grabbed/seller/target countries
It requires the transaction data, each row contains

- buyer country
- seller country
- hectare

and optionally accepts a country name or transaction id as input to show a filtered view.
The data is encoded using

- Spatial encoding - geographic location
- Color encoding left map - buyer countries (hue)
- Color encoding left map - hectares of land bought by country (luminance/saturation)
- Color encoding right map - seller countries (hue)
- Color encoding right map - hectares of land sold by country (luminance/saturation)

When starting the visualization the first thing that appears is is an overview of all countries, the left one showing the countries that grab land and the right one showing the countries that had their land grabbed.

4.3.1 Interaction

Mouseover: Hovering over a country shows the country name, the sold and bought amount of hectare. Not hovering over particular country shows the total of grabbed land in hectares, either for all countries or for the currently selected countries (if a filter is active).
Click: After clicking on a country the country filter is activated. That means that the left map shows all the countries that bought land from it, and the right map shows all the countries that the selected countries bought land from. This filter is also applied to other views.
Passive: Receiving a filter (either a single country or a transaction id) updates the maps to apply color encoding to only the relevant countries. It considers only the marked countries.
Implementation: D3 with datamaps was used to generate the maps.

4.3.2 What data does it encode?

The landgrabbing-related data to each country is obviously what we care about the most.
It might be a good idea to let the user decide what the color encoding should be used for.

Choice	Left map	Right map
#1 (default)	Amount of land grabbed by a country in ha	Amount of land grabbed from a country in ha
#2	Amount of money spent by a country for LG	Amount of money received by a country for LG
#3	Average price/ha paid by a country	Average price/ha received by a country

Only the first choice can be realized due to lack of financial data in the CSV file.

4.3.3 How does it encode the data?

It was pointed out multiple times during the lecture that the colors red and green are often bad choices because colorblind people often have difficulties distinguishing between them. However, in this particular case we have two clearly spatially distinguishable maps which could both use the same color encoding without affecting the distinguishability of the grabber countries and the grabbed countries.

- Hue
- No meaning within View 2 in particular, it just support the visual distinguishment of the left and the right map.
- Color choice: Green and Orange, should be color-blind safe according to ColorBrewer.
- Saturation
- Used for encoding the “intensity” of the quantity data. It gives an impression of the intensity of grabbing activities.
- For comparing similar countries the text values that come with the mouseover have to be used.

Using ColorBrewer the following 6-order model seems reasonable:

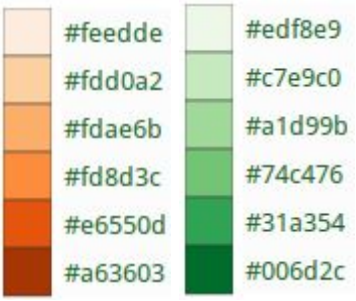

The lowest value and the highest value are taken and separated into 6 parts.

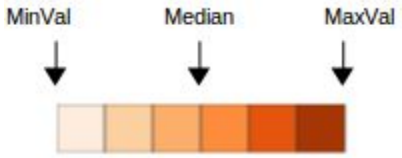

When applying this to the map we stumble upon a problem. Small countries with a high intensity are very important but hard to see:


In this example Qatar is very important but barely noticeable.

However, making the map darker and inverting the colors (high luminance = high intensity) makes it much easier to see:

In general, by using this encoding the differences are VERY clear and the high-intensity grabbers/grabbed countries are easily distinguishable.

4.3.4 Background Color

In order to encode quantifiable data with the maximum available dynamic (max. difference between lowest and highest value) the background should not be too bright and should not be too dark. It should also not be too poppy and distracting. A medium-grey color seems to be good to achieve that goal. Medium grey seems to be a reasonable choice. The irrelevant countries (not involved in landgrabbing or selection) should merge with the background. A darker shade of grey should work well for that.

4.3.5 Colorblind check

The selection color used to select particular countries is a light shade of blue. The colors we are using are green, orange, and blue.
On http://www.color-blindness.com/coblis-color-blindness-simulator/ we can check if they work for colorblind people:
Original view:

Anomalous Trichromacy works fine:

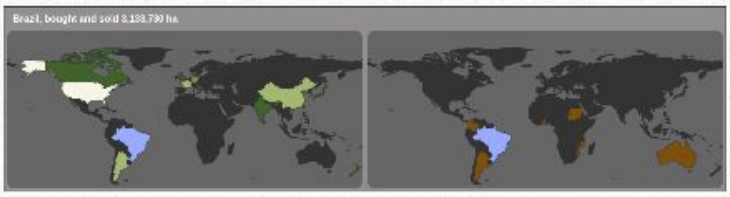

In the Dichromatic view both maps seem to use the same encoding.

But both sides are easy to distinguish because they are spatially separated, so that's okay.

4.3.6 The map size problem

If the map is too small to select a particular country then that country can be selected in View 1. Some Islands or countries that are too small to appear on the map can be selected in View 1.

4.3.7 Show information

The colored map should merely provide an intuitive overview therefore the color-encoded data should be available in text form.
Placement options:

- Popup: The data is shown in a hovering popup over the map
- Pro: Takes no space
- Con: Viewing total numbers is not possible
- Middle
- Pro: Large view, could be enhanced visually
- Con: Makes the comparison between the two maps more difficult
- Con: Takes a lot of space

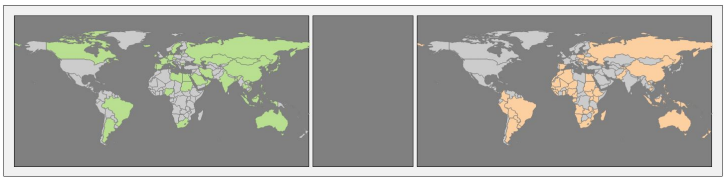

- Top
- Pro: Takes little space, not distracting

4.3.8 Problems and challenges

1. The country names used in the CSV and the country names in the datamaps json file were different and had to be “synchronized”.
2. Due to the limited size of the map the colors and borders had to be optimized in order to make it possible for small countries. This required a lot of thought and trial-and-error until the final result.
3. Three or four small islands are not visible and are not included. They can still be seen and selected in View 1.
4. The maps have to be accessible to colorblind people. The colors were copied from ColorBrewer2.org. Adjusting it properly required a lot of work.
5. There is not enough data to include financial aspects into the view.

4.3.9 Results

Global view

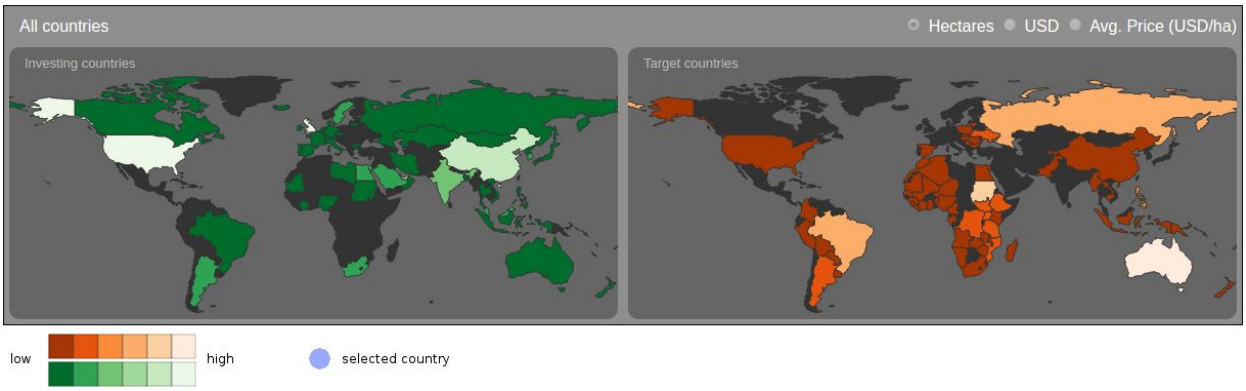

At first glance the strengths of this visualization become obvious: The biggest grabber countries in this dataset are USA, UK, China, UAE. The most grabbed country is Australia, but Russia, Brazil and Sudan also sold a lot of land.
Filter: Only Argentina

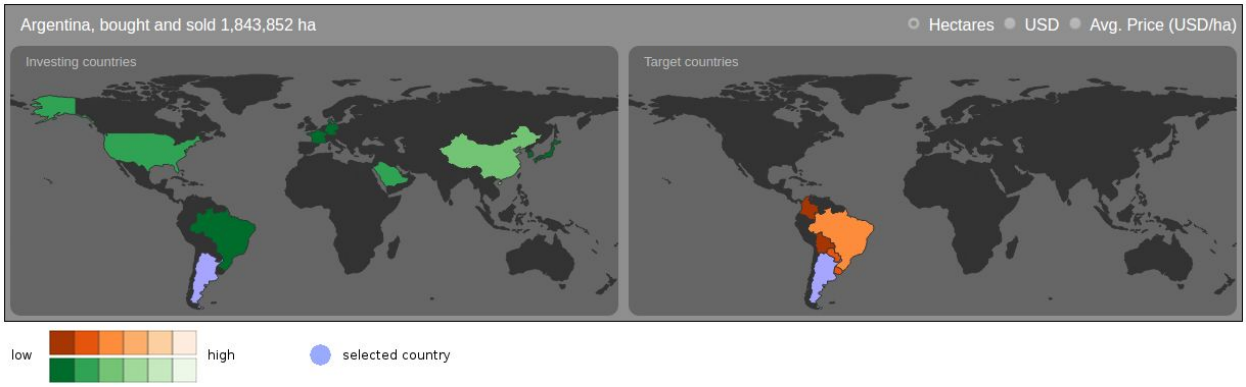

China, USA and Saudi Arabia seem to invest into Argentina while Argentina mostly grabs land from south american countries, in particular Brazil, Paraguay, Uruguay.

4.4 View 3

This part of the dashboard provides more detailed information on transactions made by countries.
How the transaction box was designed in the mock up of M2
The overall implementation of the transaction info box does not differ much from the initial mock up prototype from M2. The mock up was designed to give further and more detailed information on transactions. The box was supposed to pop up as soon as a corresponding event has been fired in the dashboard. The info box should be capable of displaying data of either a country or a transaction. To achieve this, the mock up describes two distinct views.
One shows information of a specific country, where data like production interests, economic sectors are visualized by bar charts and also contains a list of all incoming and outgoing deals of the country.
The second view is supposed to give detailed information on a particular transaction. The uppermost section gives insights on the company name, investment, production and how many hectares have been grabbed. Below, two pie charts were intended to draw the distribution of hectares in deals of the company and the country, to give users an instant overview of how many hectares have been sold or purchased in this transaction. Data like sectors, production purposes and comparisons are visualized with bar charts.
At the bottom there is a text box that provides the lore of the transaction to give interested users better understandings on the background story.
Presentation Feedback and Changes in M3
At the beginning we examined different attempts on how to properly display data in the transaction box. Although radial charts have issues in accurately presenting information, we found it to be a good idea in this case, to use pie charts to draw the amount of grabbed hectares for every transaction in order to give users a quick overview.
In the talk after our M2 presentation we received feedback, that pie charts could be converted into bar charts, since they are much better for comparison and can be perceived more accurately.

This was one reason to exclude this chart type, however, we were convinced to replace pie charts by bar charts, after we had tested them on our own. We roughly implemented them to show the grabbed hectares and as well created a bar chart as comparison, visualizing the same data. The pie charts had the expected effect: they gave a rough idea on how many hectares have been grabbed by companies and countries. However, the bar chart counterpart was able to show the information more precisely and was easier to understand at the same time.
At this point we decided to drop our initial ideas on using pie charts and decided to use bar charts instead.

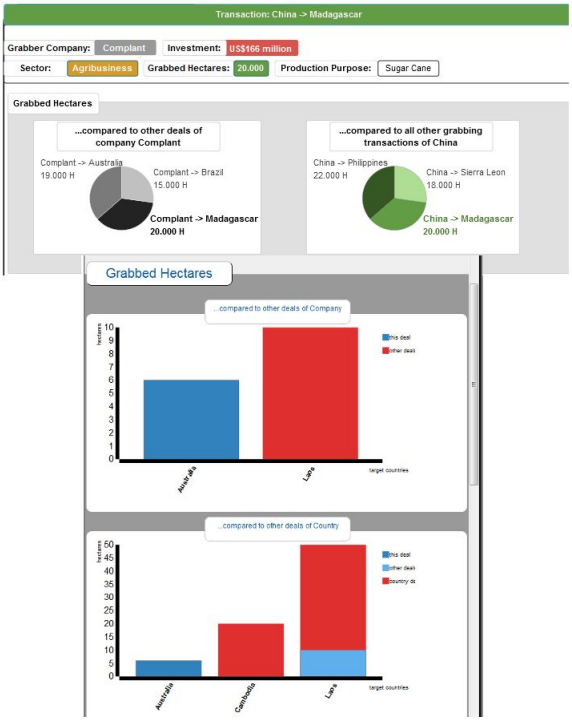

Concerning the transaction info box, this was the most major change from M2 to M3. The layout and design defined by the mock up was applied to the implementation in html and d3 without experiencing any considerable modifications.
To satisfy the demands of colorblind users the greenish color used in the mock up was transferred to a colorblind safe blueish color in the M3 implementation.
After creating the bar charts as they are specified by the prototype, text labels were unreadable, because they overlapped. Furthermore the bars were extremely small and hard to read, due to the modestly chosen bounds of the svg. To fix this issue, we increased the width and height of all charts, which caused the side effect, that charts are now aligned horizontally rather than vertically, to fit in the transaction box' bounds.

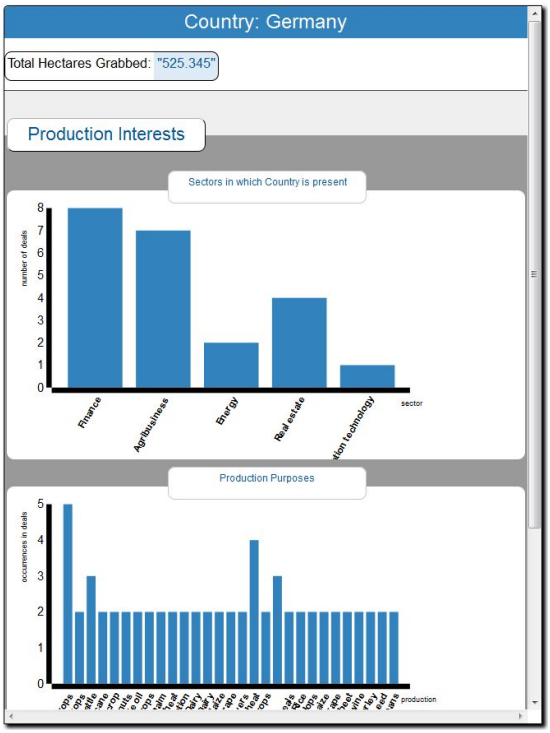

Functionality of the M3 implementation

As mentioned above, the implementation of M3 hardly deviates from the initial mock up. Although the pie charts in the transaction view have been converted into bar charts, the functionality was not affected at all. The d3 implementation creates html output depending on whether information on a country or a transaction is needed. Therefore two main methods can be used by other charts of the dashboard to generate the desired content by passing the country name or the line in the csv.

4.5 Dashboard and Linking

We have 4 views. Each of our views shows countries and/or transactions in some form or shape. If the user clicks on a country or transactions within a view with a cursor then this interaction is applied to all of the views.

This selection can contain the following information:

- a country name, represented by a string

or

- a transaction, represented by a line number from the CSV file (or indirectly, buyer country and seller country)

Each of the views has at least one create*(...) function. If it is visible from the beginning it needs no parameters. If it is created after interaction then it needs two create functions, one for country selections and one for transaction selections. It also needs a delete function.
Each few that is visible from the beginning and changes has two update*(...) function, one for country selections, one for transaction selections.
For HBC, Worldmap, Scatterplot we need

- create*();
- update*(countryname);
- update*(linenumber);

Transaction Info Box

- createTransactionInfoBox(countryname);
- createTransactionInfoBox(linenumber);
- deleteTransactionInfoBox();

When the user clicks on a country or transaction element in a view, the updateDashboard(filter); is called, a function that updates all of the views by calling all of the update functions.

5 Results

Our result is a 4-view visualization dashboard.
Now, how can various tasks be fulfilled by our tool and how does it make users capable of acquiring desired information? Since every view realizes a different chart type, the final dashboard provides great mechanics for just exploring the data and eventually stumble over an interesting fact or story. To address a user task that was specified at the beginning of the document:
"*Does Austria compete in Landgrabbing?"*
Our tool's answer is: "No". According to the data our tool is built upon, there have been no transactions made confirming incoming or outgoing purchases in Austria, thus it does not appear in the Bundle Chart as selectable item.
This is a possible task that can be solved with this tool, however it is capable of rendering a great amount of information in one dashboard.
For demonstration purposes we picked the country *China:*

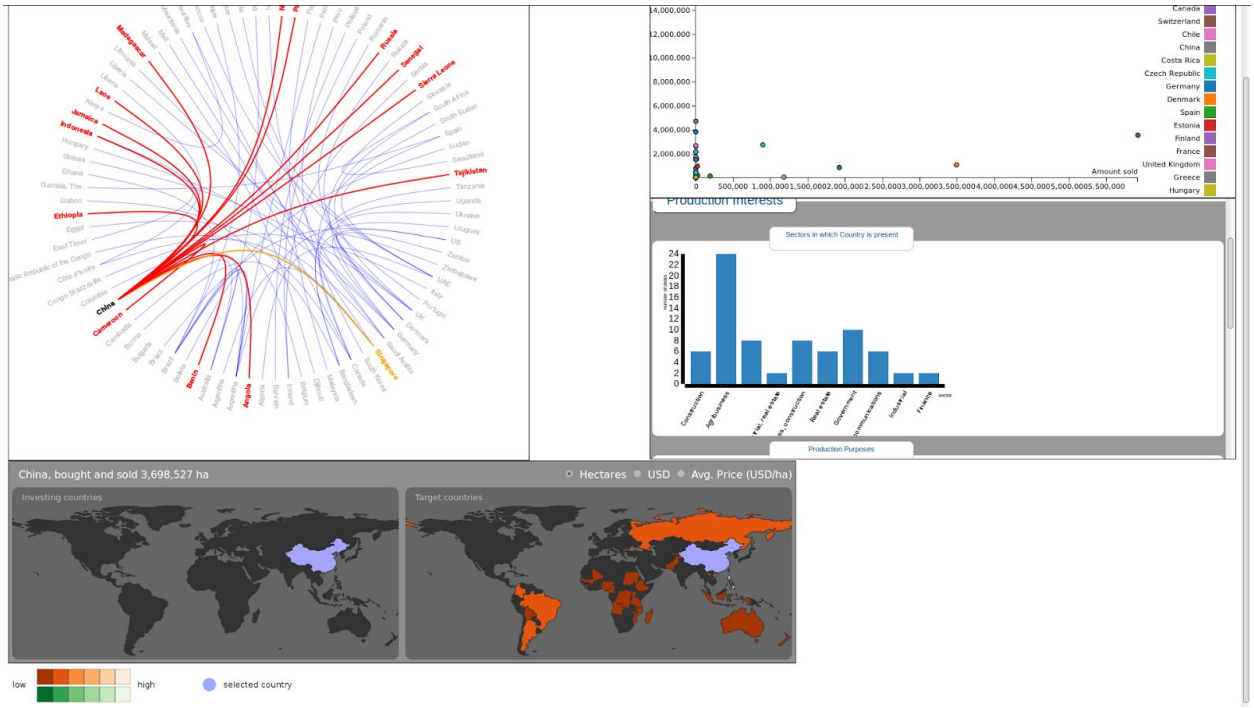

In this particular view China is selected. We can see

- the countries that were involved in landgrabbing deals (with China)
- the buyer countries in yellow
- the seller countries in red
- The geographical position of the grabber and seller countries

- the GDP/Landgrabbing scatterplot
- all the involved companies grabbing land from this country
- all the involved companies grabbing land from other countries
- production purposes
- production interests

Another example:

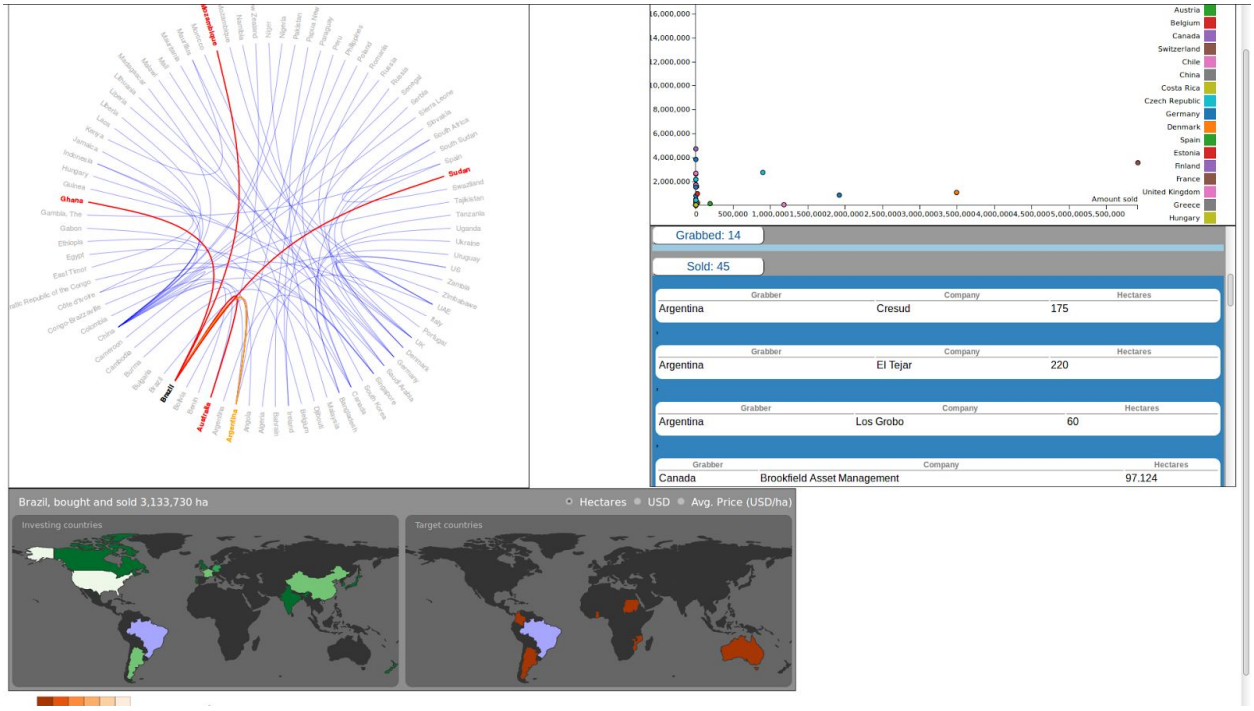

6 Discussion

6.1 Strengths

The strength of our approach and implementation clearly lies in the incorporation of the two most important realms in the data: geographical information and information concerning the different transactions. The hierarchical edge bundling chart shows all transactions available in the dataset, while the world mini maps clearly delineate the geographical origin and target of the transactions.

In addition, this is done without too much visual cluttering or hairball effects. Especially when comparing our approach to the different visualisations already in existence (see Chapter 2: Related Work) the strength of combining geographical and transaction-related information without too much visual cluttering becomes apparent.
In addition, our visualisation provides detailed information about the nature of deals (e.g. economic sector of companies, production purposes) and about the economic situation in investing / target countries. This additional level of information makes it possible to investigate the particular role of certain countries in international land investments. All this can be done within one single dashboard.

Therefore, we want to argue that, while it is possible to orient oneself easily within the different dashboard views, we have achieved a considerable depth of provided information.

6.2 Weaknesses

A possible weakness of our approach and implementation lies in the amount of data incorporated in the visualisation; available data sets on the topic of landgrabbing are scarce and ill-matched. A solution would have been to embark on major data wrangling; a task that would have certainly exceeded the scope of this design study project.

A point to take away from this issue is that a well exercised visualisation project needs a good amount of well chosen and well cleaned data.

Gross domestic product and human development index informations are not really reliable on many countries, because as mentioned it is clearly a hard task to find and merge data in terms of completeness.

7 Work separation

7.1 General responsibility

… for the design, implementation and documentation of the following vis parts:

Project Part	Member
Visualization Design General	Everyone
View 1 (Transaction View, HBC)	Aichinger
View 2 (Geographic View, Minimaps)	Doknic
View 3 (Transaction Info Box)	Thöndel
View 4 (Scatterplot)	Dernjani

For more details see M2 and M3 Written Report on Website.

7.2 M4 participation

7.2.1 Participation table

	Aichinger	Doknic	Thöndel	Dernjani
WRG	x	x	x	x
WR#1		x		
WR#2		x		
WR#3		x		x
WR#4	x	x		
WR#5		x	x	
WR#6	x			
WR#7		x		

VSG	x	x	x	x
VS#1	x			
VS#2				
VS#3				
VS#4				

WRG = Written report general
WR#N = Written report chapter N
VSG = Visualization general
VS#N = Visualization view N
PR = Presentation

x = worked on...

7.2.2 Work time table

Aichinger		
Date	**Time**	**Working on...**
20.01.2016	02:00	Final Report (general, Chapter 4.1, Chapter 6)

Doknic		
Date	**Time**	**Working on...**
16.12.2016	04:00 h	Creating M4 Final Report Document, adding notes and structure, copy content from M3, experimenting with transaction views
30.12.2016	00:45 h	Chapter 1.2, Chapter 2
16.01.2017	2:00 h	Chapter 2.1, Chapter 3, Chapter 4, Chapter 5

Thöndel		
Date	**Time**	**Working on...**
20.01.2017	1h	Written Report (Chapter 5)

Dernjani		
Date	**Time**	**Working on...**
10.12.2017	8h	Programming Scatterplot for M3 (was not logged)
21.01.2017	2h	Presentation Strengths and Weaknesses
21.01.2017	2h	General proofreading / Chapter 3.4 - Transectionbox (Final) / 6.2 Weaknesses