Recent progress on the Hegselmann-Krause bounded confidence model

Peter Hegarty
(plus: Edvin Wedin, Anders Martinsson, Mattias Danielsson, Jimmie Ekström, Jesper Johansson and Gustav Karlsson)

Department of Mathematics, Chalmers/Gothenburg University

Tuesday, 14 July, 2015

- [WH1, 2014] E. Wedin and P. Hegarty, *The Hegselmann-Krause dynamics for continuous agents and a regular opinion function do not always lead to consensus.* To appear in IEEE Trans. Automat. Control (available online).
- [WH2, 2015] E. Wedin and P. Hegarty, *A quadratic lower bound for the convergence rate in the one-dimensional Hegselmann-Krause bounded confidence dynamics*, Discrete Comput. Geom. **53** (2015), No. 2, 478–486.
- [HW, 2014] P. Hegarty and E. Wedin, *The Hegselmann-Krause dynamics for equally spaced agents*. Preprint available at `http://arxiv.org/abs/1406.0819`
- [HMW, 2014] P. Hegarty, A. Martinsson and E. Wedin, *The Hegselmann-Krause dynamics on the circle converge*. Preprint available at `http://arxiv.org/abs/1410.7330`

- [M, 2015] A. Martinsson, *An improved energy argument for the Hegselmann-Krause model*. Preprint available at `http://arxiv.org/abs/1501.02183`
- [DEJK, 2015] M. Danielsson, J. Ekström, G. Karlsson and J. Johansson, *The Hegselmann-Krause model of opinion dynamics in one and two dimensions: Phase transitions, periodicity and other phenomena*. Bachelor's Thesis, Chalmers University of Technology (2015). Available on request.

- Finite number, n say, of agents, indexed by the integers $1, 2, \ldots, n$.

- Finite number, n say, of agents, indexed by the integers $1, 2, \ldots, n$.
- Time is discrete: $t = 0, 1, \ldots$

- Finite number, n say, of agents, indexed by the integers $1, 2, \ldots, n$.
- Time is discrete: $t = 0, 1, \ldots$
- A real number $x_i(t)$ represents the opinion of agent i at time t.

- Finite number, n say, of agents, indexed by the integers $1, 2, \ldots, n$.
- Time is discrete: $t = 0, 1, \ldots$
- A real number $x_i(t)$ represents the opinion of agent i at time t.
- There is a *confidence bound* $r > 0$, which is the same for all agents.

- Finite number, n say, of agents, indexed by the integers $1, 2, \ldots, n$.
- Time is discrete: $t = 0, 1, \ldots$
- A real number $x_i(t)$ represents the opinion of agent i at time t.
- There is a *confidence bound* $r > 0$, which is the same for all agents.
- Opinions are updated synchronously according to

$$x_i(t+1) = \frac{1}{|\mathcal{N}_i(t)|} \sum_{j \in \mathcal{N}_i(t)} x_j(t),$$

where

$$\mathcal{N}_i(t) = \{j : ||x_j(t) - x_i(t)|| \leq r\}.$$

- Finite number, n say, of agents, indexed by the integers $1, 2, \ldots, n$.
- Time is discrete: $t = 0, 1, \ldots$
- A real number $x_i(t)$ represents the opinion of agent i at time t.
- There is a *confidence bound* $r > 0$, which is the same for all agents.
- Opinions are updated synchronously according to

$$x_i(t+1) = \frac{1}{|\mathcal{N}_i(t)|} \sum_{j \in \mathcal{N}_i(t)} x_j(t),$$

  where

$$\mathcal{N}_i(t) = \{j : ||x_j(t) - x_i(t)|| \leq r\}.$$

- The dynamics are unaffected by rescaling (update rule is linear), so WLOG $r = 1$.

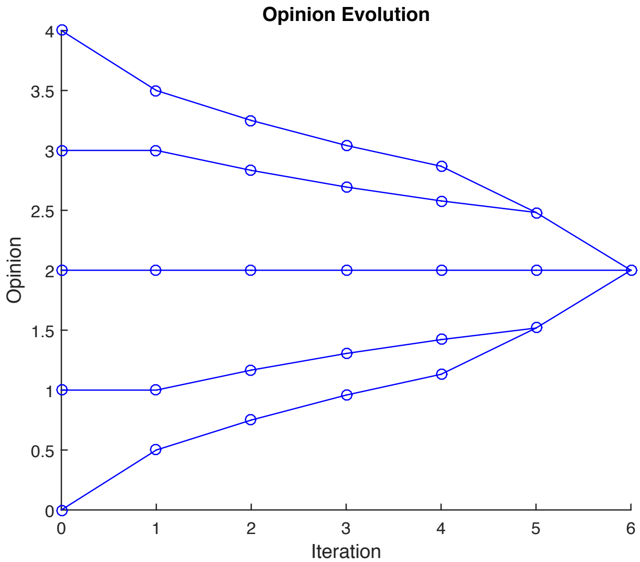
Opinion Evolution
Opinion
Iteration
4
3.5
3
2.5
2
1.5
1
0.5
0
0
1
2
3
4
5
6

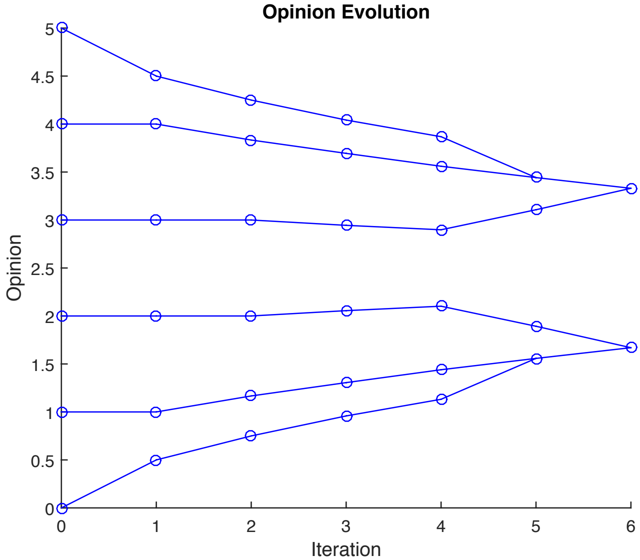
Opinion Evolution
Opinion
Iteration
0
0.5
1
1.5
2
2.5
3
3.5
4
4.5
5
0
1
2
3
4
5
6

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

Example 1. Higher dimensional Euclidean space $V = \mathbb{R}^k$.

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

Example 1. Higher dimensional Euclidean space $V = \mathbb{R}^k$.

Interpretation: There are k issues, and two agents must be close on all issues for compromise to occur.

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

Example 1. Higher dimensional Euclidean space $V = \mathbb{R}^k$.

Interpretation: There are k issues, and two agents must be close on all issues for compromise to occur. Note that a priori no reason to favour L^2-norm over any other in this interpretation.

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

Example 1. Higher dimensional Euclidean space $V = \mathbb{R}^k$.

Interpretation: There are k issues, and two agents must be close on all issues for compromise to occur. Note that a priori no reason to favour L^2-norm over any other in this interpretation. However, it gives the most natural geometrical interpretation, and is the one used in robotics applications (multi-agent rendezvous).

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

Example 1. Higher dimensional Euclidean space $V = \mathbb{R}^k$.

Interpretation: There are k issues, and two agents must be close on all issues for compromise to occur. Note that a priori no reason to favour L^2-norm over any other in this interpretation. However, it gives the most natural geometrical interpretation, and is the one used in robotics applications (multi-agent rendezvous).

Example 2. The circle $V = \mathbb{T}^1$, of diameter greater than 2.

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

Example 1. Higher dimensional Euclidean space $V = \mathbb{R}^k$.

Interpretation: There are k issues, and two agents must be close on all issues for compromise to occur. Note that a priori no reason to favour L^2-norm over any other in this interpretation. However, it gives the most natural geometrical interpretation, and is the one used in robotics applications (multi-agent rendezvous).

Example 2. The circle $V = \mathbb{T}^1$, of diameter greater than 2.

Interpretation: Imagine, for example, that the issue under discussion is the time of day or year for holding some event.

Convergence in $\mathbb{R}$:

Convergence in $\mathbb{R}$:

- Very easy to show that opinions converge to limiting values (general nonsense, Banach Fixed Point Theorem blah blah ...)

Convergence in $\mathbb{R}$:

- Very easy to show that opinions converge to limiting values (general nonsense, Banach Fixed Point Theorem blah blah ...)
- In fact quite easy to show that opinions **freeze**, i.e.: there is always some $T > 0$ such that $x_i(t) = x_i(T)$ for all i and all $t \geq T$.

Convergence in $\mathbb{R}$:

- Very easy to show that opinions converge to limiting values (general nonsense, Banach Fixed Point Theorem blah blah ...)
- In fact quite easy to show that opinions **freeze**, i.e.: there is always some $T > 0$ such that $x_i(t) = x_i(T)$ for all i and all $t \geq T$.
- Still quite easy to show that the freezing time is bounded by a universal polynomial function of the number of agents:

Convergence in $\mathbb{R}$:

- Very easy to show that opinions converge to limiting values (general nonsense, Banach Fixed Point Theorem blah blah ...)
- In fact quite easy to show that opinions **freeze**, i.e.: there is always some $T > 0$ such that $x_i(t) = x_i(T)$ for all i and all $t \geq T$.
- Still quite easy to show that the freezing time is bounded by a universal polynomial function of the number of agents: $\Rightarrow$ Can get a bound of around $O(n^5)$ from general theory of Markov chains on graphs.

Convergence in $\mathbb{R}$:

- Very easy to show that opinions converge to limiting values (general nonsense, Banach Fixed Point Theorem blah blah ...)
- In fact quite easy to show that opinions **freeze**, i.e.: there is always some $T > 0$ such that $x_i(t) = x_i(T)$ for all i and all $t \geq T$.
- Still quite easy to show that the freezing time is bounded by a universal polynomial function of the number of agents:
  $\Rightarrow$ Can get a bound of around $O(n^5)$ from general theory of Markov chains on graphs.
  $\Rightarrow$ Best to date is $O(n^3)$. Elementary argument which considers the behaviour of the extremal agents (Bhattachrya et al, 2013).

Lower bounds on a universal freezing time first studied in any seriousness by Wedin and myself [WH2, HW 2014].

Lower bounds on a universal freezing time first studied in any seriousness by Wedin and myself [WH2, HW 2014].

- Easy to see that n agents placed distance one apart will take time $\Omega(n)$ to freeze.

Lower bounds on a universal freezing time first studied in any seriousness by Wedin and myself [WH2, HW 2014].

- Easy to see that n agents placed distance one apart will take time $\Omega(n)$ to freeze.
  We proved that this configuration evolves periodically, with groups of 3 agents breaking loose at each end every 5th time step.

Lower bounds on a universal freezing time first studied in any seriousness by Wedin and myself [WH2, HW 2014].

- Easy to see that n agents placed distance one apart will take time $\Omega(n)$ to freeze.
  We proved that this configuration evolves periodically, with groups of 3 agents breaking loose at each end every 5th time step.
  In particular, the freezing time is $5n/6 + O(1)$.

Lower bounds on a universal freezing time first studied in any seriousness by Wedin and myself [WH2, HW 2014].

- Easy to see that n agents placed distance one apart will take time $\Omega(n)$ to freeze.
  We proved that this configuration evolves periodically, with groups of 3 agents breaking loose at each end every 5th time step.
  In particular, the freezing time is $5n/6 + O(1)$.
- We were surprised to discover a configuration which takes time $\Omega(n^2)$ to freeze: **Dumbbell graph**

Lower bounds on a universal freezing time first studied in any seriousness by Wedin and myself [WH2, HW 2014].

- Easy to see that n agents placed distance one apart will take time $\Omega(n)$ to freeze.
  We proved that this configuration evolves periodically, with groups of 3 agents breaking loose at each end every 5th time step.
  In particular, the freezing time is $5n/6 + O(1)$.
- We were surprised to discover a configuration which takes time $\Omega(n^2)$ to freeze: **Dumbbell graph**
- We believe that the freezing time is always $O(n^2)$, but this remains open.

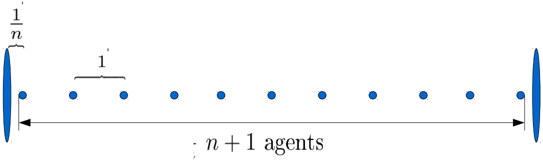

Figure: Schematic representation of the configuration $\mathcal{D}_n$. Each dumbbell has weight n.

Convergence in $\mathbb{R}^k$, $k > 1$:

Convergence in $\mathbb{R}^k$, $k > 1$:

- Opinions still freeze in finite time: this just requires a "convexity argument", which works in any Euclidean space.

Convergence in $\mathbb{R}^k$, $k > 1$:

- Opinions still freeze in finite time: this just requires a "convexity argument", which works in any Euclidean space.
- Method of Bhattachrya et al does not generalise.

Convergence in $\mathbb{R}^k$, $k > 1$:

- Opinions still freeze in finite time: this just requires a "convexity argument", which works in any Euclidean space.
- Method of Bhattachrya et al does not generalise.
  Moral: The jump from one to two dimensions is where all the new action lies.

Convergence in $\mathbb{R}^k$, $k > 1$:

- Opinions still freeze in finite time: this just requires a "convexity argument", which works in any Euclidean space.
- Method of Bhattachrya et al does not generalise.
  Moral: The jump from one to two dimensions is where all the new action lies.
- Instead, state-of-the-art for $k > 1$ is an **energy reduction argument**.

Convergence in $\mathbb{R}^k$, $k > 1$:

- Opinions still freeze in finite time: this just requires a "convexity argument", which works in any Euclidean space.
- Method of Bhattachrya et al does not generalise.
  Moral: The jump from one to two dimensions is where all the new action lies.
- Instead, state-of-the-art for $k > 1$ is an **energy reduction argument**.
  The **energy** of a Hegselmann-Krause system $\mathbf{x} = (x_1, \ldots, x_n)$ is given by

$$\mathcal{E}(\mathbf{x}) = \sum_{i,j=1}^{n} \max\{1, ||x_i - x_j||^2\}.$$

Convergence in $\mathbb{R}^k$, $k > 1$:

- Opinions still freeze in finite time: this just requires a "convexity argument", which works in any Euclidean space.
- Method of Bhattachrya et al does not generalise.
  Moral: The jump from one to two dimensions is where all the new action lies.
- Instead, state-of-the-art for $k > 1$ is an **energy reduction argument**.
  The **energy** of a Hegselmann-Krause system $\mathbf{x} = (x_1, \dots, x_n)$ is given by

$$\mathcal{E}(\mathbf{x}) = \sum_{i,j=1}^{n} \max\{1, ||x_i - x_j||^2\}.$$

  Basic Result: The dynamics always decrease the energy.

$$\mathcal{E}(\mathbf{x}(t)) - \mathcal{E}(\mathbf{x}(t+1)) \geq 4 \cdot \sum^{n} ||x_i(t) - x_i(t+1)||^2.$$

Convergence in $\mathbb{R}^k$, continued:

Convergence in $\mathbb{R}^k$, continued:

- Martinsson [M, 2015] proved a bound of $O(n^4)$, valid in all dimensions.

Convergence in $\mathbb{R}^k$, continued:

- Martinsson [M, 2015] proved a bound of $O(n^4)$, valid in all dimensions.
  His result is in a sense best-possible. No better bound can be obtained using just the energy reduction technique.

Convergence in $\mathbb{R}^k$, continued:

- Martinsson [M, 2015] proved a bound of $O(n^4)$, valid in all dimensions.
  His result is in a sense best-possible. No better bound can be obtained using just the energy reduction technique.
- **N.B.** The above only works for the L^2-norm.

Convergence in $\mathbb{R}^k$, continued:

- Martinsson [M, 2015] proved a bound of $O(n^4)$, valid in all dimensions.
  His result is in a sense best-possible. No better bound can be obtained using just the energy reduction technique.
- **N.B.** The above only works for the L^2-norm.
- For lower bounds, n agents placed equidistantly around a circle will also require time $\Omega(n^2)$ to freeze.

Convergence in $\mathbb{R}^k$, continued:

- Martinsson [M, 2015] proved a bound of $O(n^4)$, valid in all dimensions.
  His result is in a sense best-possible. No better bound can be obtained using just the energy reduction technique.
- **N.B.** The above only works for the L^2-norm.
- For lower bounds, n agents placed equidistantly around a circle will also require time $\Omega(n^2)$ to freeze.
  This is a genuinely 2-dimensional example. Also, in contrast to the dumbbell, this configuration reaches consensus.

Convergence in $\mathbb{R}^k$, continued:

- Martinsson [M, 2015] proved a bound of $O(n^4)$, valid in all dimensions.
  His result is in a sense best-possible. No better bound can be obtained using just the energy reduction technique.
- **N.B.** The above only works for the L^2-norm.
- For lower bounds, n agents placed equidistantly around a circle will also require time $\Omega(n^2)$ to freeze.
  This is a genuinely 2-dimensional example. Also, in contrast to the dumbbell, this configuration reaches consensus.
- We believe that the freezing time is $O(n^2)$ in all dimensions.

Convergence on $\mathbb{T}^1$:

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one.
  However, there are even non-periodic frozen configurations.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one.
  However, there are even non-periodic frozen configurations.
- Hendrickx et al (2009) asked if opinions must always converge on the circle.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one.
  However, there are even non-periodic frozen configurations.
- Hendrickx et al (2009) asked if opinions must always converge on the circle.
  Proven by us [HMW, 2014].

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one.
  However, there are even non-periodic frozen configurations.
- Hendrickx et al (2009) asked if opinions must always converge on the circle.
  Proven by us [HMW, 2014]. Proof uses both the energy reduction technqiue and a modification of the idea in Bhattachrya et al, both suitably modified for the circle.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one.
  However, there are even non-periodic frozen configurations.
- Hendrickx et al (2009) asked if opinions must always converge on the circle.
  Proven by us [HMW, 2014]. Proof uses both the energy reduction technqiue and a modification of the idea in Bhattachrya et al, both suitably modified for the circle.
- The influence digraph can change at most $O(n^4)$ times.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one.
  However, there are even non-periodic frozen configurations.
- Hendrickx et al (2009) asked if opinions must always converge on the circle.
  Proven by us [HMW, 2014]. Proof uses both the energy reduction technqiue and a modification of the idea in Bhattachrya et al, both suitably modified for the circle.
- The influence digraph can change at most $O(n^4)$ times. However, it can take arbitrarily long for these changes to occur.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one. However, there are even non-periodic frozen configurations.
- Hendrickx et al (2009) asked if opinions must always converge on the circle.
  Proven by us [HMW, 2014]. Proof uses both the energy reduction technqiue and a modification of the idea in Bhattachrya et al, both suitably modified for the circle.
- The influence digraph can change at most $O(n^4)$ times. However, it can take arbitrarily long for these changes to occur.
- Remains open to prove convergence in $\mathbb{T}^k$.

Much less is known here. There is an absence of basic techniques to deal with the complex dependencies between agents which arise in the dynamics, even if initial opinions are drawn **independently** from some probability distribution.

Much less is known here. There is an absence of basic techniques to deal with the complex dependencies between agents which arise in the dynamics, even if initial opinions are drawn **independently** from some probability distribution.

However, we can appeal to two basic principles which together would lead us to expect typical **phase transition** behaviour.

Much less is known here. There is an absence of basic techniques to deal with the complex dependencies between agents which arise in the dynamics, even if initial opinions are drawn **independently** from some probability distribution.

However, we can appeal to two basic principles which together would lead us to expect typical **phase transition** behaviour.

- **Monotonicity:**

Much less is known here. There is an absence of basic techniques to deal with the complex dependencies between agents which arise in the dynamics, even if initial opinions are drawn **independently** from some probability distribution.

However, we can appeal to two basic principles which together would lead us to expect typical **phase transition** behaviour.

- **Monotonicity:** Dilating the opinion space without changing the "relative distribution" of opinions should always make consensus less likely.

Much less is known here. There is an absence of basic techniques to deal with the complex dependencies between agents which arise in the dynamics, even if initial opinions are drawn **independently** from some probability distribution.
However, we can appeal to two basic principles which together would lead us to expect typical **phase transition** behaviour.

- **Monotonicity:** Dilating the opinion space without changing the "relative distribution" of opinions should always make consensus less likely.
- **Zero-One Law:**

Much less is known here. There is an absence of basic techniques to deal with the complex dependencies between agents which arise in the dynamics, even if initial opinions are drawn **independently** from some probability distribution.

However, we can appeal to two basic principles which together would lead us to expect typical **phase transition** behaviour.

- **Monotonicity:** Dilating the opinion space without changing the "relative distribution" of opinions should always make consensus less likely.
- **Zero-One Law:** Suppose initial opinions are chosen *independently* from some fixed distribution. As $n \to \infty$, the probability of reaching consensus should go to 0 or 1, i.e.: there should be a "typical behaviour".

Much less is known here. There is an absence of basic techniques to deal with the complex dependencies between agents which arise in the dynamics, even if initial opinions are drawn **independently** from some probability distribution.

However, we can appeal to two basic principles which together would lead us to expect typical **phase transition** behaviour.

- **Monotonicity:** Dilating the opinion space without changing the "relative distribution" of opinions should always make consensus less likely.
- **Zero-One Law:** Suppose initial opinions are chosen *independently* from some fixed distribution. As $n \to \infty$, the probability of reaching consensus should go to 0 or 1, i.e.: there should be a "typical behaviour".

Monotonicity seems intuitively obvious, but the $0 - 1$ principle perhaps need more motivation. Note that **nothing** is proven however.

$0-1$ Law and the Continuous Agent Model (CAM):

$0-1$ Law and the Continuous Agent Model (CAM):

- **Basic Idea:** Instead of drawing opinions independently from a (continuous) distribution $f(x)$, consider a continuum of agents with $f(x)$ describing an opinion density function.

$0-1$ Law and the Continuous Agent Model (CAM):

- **Basic Idea:** Instead of drawing opinions independently from a (continuous) distribution $f(x)$, consider a continuum of agents with $f(x)$ describing an opinion density function.
- **Simplest example:** Uniformly independent opinions on $[0, L]$ corresponds to the opinion function $x_0 : [0, 1] \rightarrow \mathbb{R}$, $x_0(\alpha) = L\alpha$.

$0-1$ Law and the Continuous Agent Model (CAM):

- **Basic Idea:** Instead of drawing opinions independently from a (continuous) distribution $f(x)$, consider a continuum of agents with $f(x)$ describing an opinion density function.
- **Simplest example:** Uniformly independent opinions on $[0, L]$ corresponds to the opinion function $x_0 : [0, 1] \to \mathbb{R}$, $x_0(\alpha) = L\alpha$.
- **The dynamics:**

$$x_{t+1}(\alpha) = \frac{1}{\mu(\mathcal{N}_t(\alpha))} \int_{\mathcal{N}_t(\alpha)} x_t(\beta)\, d\beta,$$

where $\mathcal{N}_t(\alpha) = \{\beta : ||x_t(\beta) - x_t(\alpha)|| \leq 1\}$ and μ is Lebesgue measure.

$0-1$ Law and the Continuous Agent Model (CAM):

- **Basic Idea:** Instead of drawing opinions independently from a (continuous) distribution $f(x)$, consider a continuum of agents with $f(x)$ describing an opinion density function.
- **Simplest example:** Uniformly independent opinions on $[0, L]$ corresponds to the opinion function $x_0 : [0, 1] \to \mathbb{R}$, $x_0(\alpha) = L\alpha$.
- **The dynamics:**

$$x_{t+1}(\alpha) = \frac{1}{\mu(\mathcal{N}_t(\alpha))} \int_{\mathcal{N}_t(\alpha)} x_t(\beta)\, d\beta,$$

  where $\mathcal{N}_t(\alpha) = \{\beta : ||x_t(\beta) - x_t(\alpha)|| \leq 1\}$ and μ is Lebesgue measure.
- **Precise formulation of $0-1$ Law:** For independent initial opinions, as $n \to \infty$ one almost surely reaches consensus if and only if one reaches consensus in CAM.

- I regard it as the main open problem about the HK-model, to establish the relationship between CAM and large random configurations.

- I regard it as the main open problem about the HK-model, to establish the relationship between CAM and large random configurations. The matter seems to be very subtle, there are no proofs of things which might seem "obvious".

- I regard it as the main open problem about the HK-model, to establish the relationship between CAM and large random configurations. The matter seems to be very subtle, there are no proofs of things which might seem "obvious".
- It is not known in general if a configuration of opinions in CAM always converges to something (Hendrickx et al, 2009).

- I regard it as the main open problem about the HK-model, to establish the relationship between CAM and large random configurations. The matter seems to be very subtle, there are no proofs of things which might seem “obvious”.
- It is not known in general if a configuration of opinions in CAM always converges to something (Hendrickx et al, 2009).
- Wedin and I [WH1, 2014] gave the first example of a **regular** opinion function (piecewise differentiable, with positive lower and upper bounds on the derivative) which never reaches consensus.

- I regard it as the main open problem about the HK-model, to establish the relationship between CAM and large random configurations. The matter seems to be very subtle, there are no proofs of things which might seem "obvious".
- It is not known in general if a configuration of opinions in CAM always converges to something (Hendrickx et al, 2009).
- Wedin and I [WH1, 2014] gave the first example of a **regular** opinion function (piecewise differentiable, with positive lower and upper bounds on the derivative) which never reaches consensus. Even this is a non-trivial task.

- I regard it as the main open problem about the HK-model, to establish the relationship between CAM and large random configurations. The matter seems to be very subtle, there are no proofs of things which might seem “obvious”.
- It is not known in general if a configuration of opinions in CAM always converges to something (Hendrickx et al, 2009).
- Wedin and I [WH1, 2014] gave the first example of a **regular** opinion function (piecewise differentiable, with positive lower and upper bounds on the derivative) which never reaches consensus. Even this is a non-trivial task.
- Our example is a kind of **double-S**.

- I regard it as the main open problem about the HK-model, to establish the relationship between CAM and large random configurations. The matter seems to be very subtle, there are no proofs of things which might seem "obvious".
- It is not known in general if a configuration of opinions in CAM always converges to something (Hendrickx et al, 2009).
- Wedin and I [WH1, 2014] gave the first example of a **regular** opinion function (piecewise differentiable, with positive lower and upper bounds on the derivative) which never reaches consensus. Even this is a non-trivial task.
- Our example is a kind of **double-S**.
- Problem remains open for linear functions (those corresponding to a uniform distribution of opinions).

Simulations:

Simulations:

- Many simulations performed in [DEJK, 2015] for uniform distributions of agents in regions of $\mathbb{R}^1$ and $\mathbb{R}^2$.

Simulations:

- Many simulations performed in [DEJK, 2015] for uniform distributions of agents in regions of $\mathbb{R}^1$ and $\mathbb{R}^2$.
- In $\mathbb{R}^1$ there is only one "region", namely an interval.

Simulations:

- Many simulations performed in [DEJK, 2015] for uniform distributions of agents in regions of $\mathbb{R}^1$ and $\mathbb{R}^2$.
- In $\mathbb{R}^1$ there is only one "region", namely an interval. Simulations give overwhelming evidence for existence of a critical length, slightly above 5.

Simulations:

- Many simulations performed in [DEJK, 2015] for uniform distributions of agents in regions of $\mathbb{R}^1$ and $\mathbb{R}^2$.
- In $\mathbb{R}^1$ there is only one "region", namely an interval. Simulations give overwhelming evidence for existence of a critical length, slightly above 5.

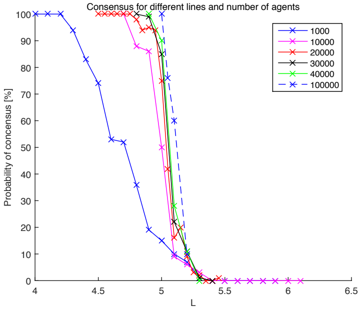

- In $\mathbb{R}^2$, the critical area will depend on the shape of the region.

- In $\mathbb{R}^2$, the critical area will depend on the shape of the region.
- Simulations seem to indicate that critical area $= +\infty$ for a disc. Not proven though.

- In $\mathbb{R}^2$, the critical area will depend on the shape of the region.
- Simulations seem to indicate that critical area $= +\infty$ for a disc. Not proven though.
- Finite critical area for a regular n-gon.

- In $\mathbb{R}^2$, the critical area will depend on the shape of the region.
- Simulations seem to indicate that critical area $= +\infty$ for a disc. Not proven though.
- Finite critical area for a regular n-gon. But more data needed to get a good idea of how the critical area depends on n.

- In $\mathbb{R}^2$, the critical area will depend on the shape of the region.
- Simulations seem to indicate that critical area $= +\infty$ for a disc. Not proven though.
- Finite critical area for a regular n-gon. But more data needed to get a good idea of how the critical area depends on n.
- So far, our simulations have not yielded any "shocking" findings, but lots of data needed in these studies.

- In $\mathbb{R}^2$, the critical area will depend on the shape of the region.
- Simulations seem to indicate that critical area $= +\infty$ for a disc. Not proven though.
- Finite critical area for a regular n-gon. But more data needed to get a good idea of how the critical area depends on n.
- So far, our simulations have not yielded any "shocking" findings, but lots of data needed in these studies. Also, average freezing times jump when one is close to a critical area, because of the tendency for **semi-stable configurations** to form.

In [DEJK, 2015] a nice link was discovered between:

In [DEJK, 2015] a nice link was discovered between:
(1) The ultimately periodic evolution of a **half-plane configuration** on $\mathbb{Z} \times \mathbb{Z}_{+}$.

In [DEJK, 2015] a nice link was discovered between:
(1) The ultimately periodic evolution of a **half-plane configuration** on $\mathbb{Z} \times \mathbb{Z}_+$.
(2) The one-dimensional HK-model with a **self-weight** $w > 1$.

In [DEJK, 2015] a nice link was discovered between:
(1) The ultimately periodic evolution of a **half-plane configuration** on $\mathbb{Z} \times \mathbb{Z}_+$.
(2) The one-dimensional HK-model with a **self-weight** $w > 1$.
Call this model HK-w. Dynamics given by

$$x_i(t+1) = \frac{1}{w + |\mathcal{N}_i^*(t)|} \left(w \cdot x_i(t) + \sum_{j \in \mathcal{N}_i^*(t)} x_j(t) \right),$$

where $\mathcal{N}_i^*(t) = \mathcal{N}_i(t) \backslash \{i\}$.

In [DEJK, 2015] a nice link was discovered between:
(1) The ultimately periodic evolution of a **half-plane configuration** on $\mathbb{Z} \times \mathbb{Z}_+$.
(2) The one-dimensional HK-model with a **self-weight** $w > 1$. Call this model HK-w. Dynamics given by

$$x_i(t+1) = \frac{1}{w + |\mathcal{N}_i^*(t)|} \left(w \cdot x_i(t) + \sum_{j \in \mathcal{N}_i^*(t)} x_j(t) \right),$$

where $\mathcal{N}_i^*(t) = \mathcal{N}_i(t) \backslash \{i\}$.
(3) The HK-model in **continuous time** (HKCT):

$$\frac{dx_i}{dt} = \sum_{j \in \mathcal{N}_i(t)} (x_j - x_i).$$

- The half-plane configuration is equivalent to HK-3.

- The half-plane configuration is equivalent to HK-3.
- Let $w \to \infty$ and simoultaneously rescale time by a factor w. Then HK-w "approaches" HKCT in some sense, but some of the details of this remain unclear.

- The half-plane configuration is equivalent to HK-3.
- Let $w \to \infty$ and simoultaneously rescale time by a factor w. Then HK-w "approaches" HKCT in some sense, but some of the details of this remain unclear.
- For example, there is not always a unique solution to HKCT. Does there always exist a unique limit to HK-w ?

- The half-plane configuration is equivalent to HK-3.
- Let $w \to \infty$ and simoultaneously rescale time by a factor w. Then HK-w "approaches" HKCT in some sense, but some of the details of this remain unclear.
- For example, there is not always a unique solution to HKCT. Does there always exist a unique limit to HK-w ?
- We can also show that, in HKCT, the influence graph always stabilises in time $O(n^2)$.

Other Open Problems:

Other Open Problems:

- **Heterogeneous HK:** Different agents are allowed to have different radii of confidence and/or influence.

Other Open Problems:

- **Heterogeneous HK:** Different agents are allowed to have different radii of confidence and/or influence. These models studied in (Mirtabatabaei-Bullo, 2012), but even the most basic issues unresolved, e.g.: do opinions always converge ?

Other Open Problems:

- **Heterogeneous HK:** Different agents are allowed to have different radii of confidence and/or influence. These models studied in (Mirtabatabaei-Bullo, 2012), but even the most basic issues unresolved, e.g.: do opinions always converge ?
- **Non-metric models in $\mathbb{R}^k$:** Suppose agents are willing to compromise if they are close on at least l issues, for some $l < k$.

Other Open Problems:

- **Heterogeneous HK:** Different agents are allowed to have different radii of confidence and/or influence. These models studied in (Mirtabatabaei-Bullo, 2012), but even the most basic issues unresolved, e.g.: do opinions always converge ?
- **Non-metric models in $\mathbb{R}^k$:** Suppose agents are willing to compromise if they are close on at least l issues, for some $l < k$.
  This seems like a natural relaxation of the HK-rule, but means that the "distance" between agents will not satisfy the triangle inequality.

Other Open Problems:

- **Heterogeneous HK:** Different agents are allowed to have different radii of confidence and/or influence. These models studied in (Mirtabatabaei-Bullo, 2012), but even the most basic issues unresolved, e.g.: do opinions always converge ?
- **Non-metric models in $\mathbb{R}^k$:** Suppose agents are willing to compromise if they are close on at least l issues, for some $l < k$.
  This seems like a natural relaxation of the HK-rule, but means that the "distance" between agents will not satisfy the triangle inequality.
  I only know of one paper (Zhang-Sun, 2009) which even mentions this idea.

Other Open Problems:

- **Heterogeneous HK:** Different agents are allowed to have different radii of confidence and/or influence. These models studied in (Mirtabatabaei-Bullo, 2012), but even the most basic issues unresolved, e.g.: do opinions always converge ?
- **Non-metric models in $\mathbb{R}^k$:** Suppose agents are willing to compromise if they are close on at least l issues, for some $l < k$.
  This seems like a natural relaxation of the HK-rule, but means that the "distance" between agents will not satisfy the triangle inequality.
  I only know of one paper (Zhang-Sun, 2009) which even mentions this idea. They just do some simulations, but no rigorous results whatsoever are in the literature.